Imperial College London

# Active Network Management for Improved Flexibility in Distribution Networks

Julio Pérez Olvera

30 September 2020

Control & Power Group
Department of Electrical and Electronic Engineering
Imperial College London

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Electrical Engineering of Imperial College London
and the Diploma of Imperial College London

# Abstract

Decarbonisation efforts have prompted the rapid uptake of low-carbon technologies (LCTs) in distribution networks. Their power use and delivery are highly variable and can cause various voltage-regulation and thermal-overload problems, thus making it difficult to maintain a secure operation of these networks. The problem is compounded by the lack of instrumentation in low-voltage networks. Two specific problems are identified: 1) the need for greater flexibility in power routing in distribution networks and 2) the need for optimising controllers operating with partial visibility of the network state and with practical computational efficiency. Power flow flexibility is provided by AC/DC power electronic converters and DC links, as part of an Active Network Management (ANM) scheme. An Optimal Power Flow (OPF) algorithm is modified to include DC links and serves as a benchmark for dispatching the DC power flows but is deemed too computationally intensive for real-time control.

Alternatives to the OPF, based on regression and reinforcement learning, were examined using case-study networks. Regression mapping of the OPF output is proposed to move the computational effort offline and to address the partial observability of the network. However, it was found to have poor accuracy and a limited ability to generalise. For better generalisation and accuracy, a model-free reinforcement learning method based on deep Q networks was proposed. The performance was on a par with the OPF while being able to generalise to different loading conditions, however the training required several hours. A model-based algorithm using an explicit model of the network was proposed to further improve the accuracy and reduce the computational burden. The training time was in the range of minutes with good accuracy compared to the OPF.

ANM using power converters together with controllers based on reinforcement learning represent a compelling alternative to further integrate LCTs in distribution networks.

# Declaration of Originality

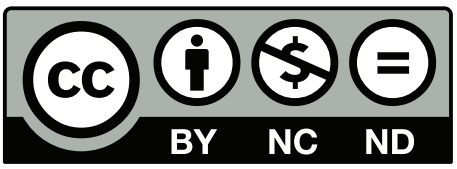
I hereby declare that the material included in this thesis is the result of my own research and work. Ideas, work and diagrams made by others are appropriately referenced. The work presented here has not been used in support for another degree or qualification in this on any other academic institution.

Julio Pérez Olvera

30 September 2020

# Copyright Statement

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND). Under this licence, you may copy and redistribute the material in any medium or format on the condition that; you credit the author, do not use it for commercial purposes and do not distribute modified versions of the work. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.



# Acknowledgements

First, I would like to thank my mother, my father and my little sister for their support and encouragement, I would not be where I am without your support and wisdom. *Gracias Bety, Ale y Diana por su apoyo incondicional.*

I would like to thank my supervisor Professor Timothy Green, for his constant support and guidance throughout these four years, for the flexibility and freedom he gave me to pursue the research areas that interest me the most. Thank you Professor Green for your guidance and advice. Also, I want to express my gratitude to my co-supervisor Dr Adrià Junyent-Ferré for his valuable advice, for the thoughtful meetings and for supporting me on my side projects. I would also like to thank Dr Michael Fairbank, at the University of Essex, for introducing me to his work on Adaptive Dynamic Programming and Neural Networks and for his willingness to collaborate during the last stage of my PhD. I am grateful to Western Power Distribution and EA Technology Ltd for granting us access to the OpenLV project data and for the technical assistance provided during this project. I would also like to take this opportunity to thank my Viva examiners, Professor Jianzhong Wu and Professor Bikash Pal, for their feedback and suggestions, which improved this thesis.

I am extremely grateful for the financial support provided by the Mexican Energy Ministry (SENER) and the National Council for Science and Technology (CONACYT) to fund my PhD studies through the PhD Scholarship number: 441647. Without their financial support I would not have been able to pursue a PhD abroad in one of the leading institutions worldwide. Also, I would like to thank Professor Guillermo Hernandez Gonzalez, Professor Ricardo Fernandez del Busto and Professor Enrique Muñoz Diaz for their early encouragement to pursue both an MSc and a PhD. Their teaching on electric circuits, industrial power systems and control during my time at ITESM CCM was one of the reasons I decided to pursue further education in this field. To Eduardo Nolasco and Esteve Borràs, my long-time friends, thank you for the very interesting conversations and for the great memories.

I would like to thank everyone in the Control & Power group at Imperial College for their friendship and support. It was great to be a part of the community! A special mention to Robert Brown at London South Bank University, with whom I've had very thoughtful discussions and have collaborated closely during the final year of my PhD. Thank you for all the suggestions and feedback Robert!

Finally, a special thank you to my partner Karolína Lukoszová, who has been there for me through the ups and downs of this long journey. She has supported me and encouraged me all along, always with a smile on her face. Thank you for everything Karolinka!

# Contents

**Abstract** 2

**Acknowledgement** 5

**Contents** 6

**List of Figures** 9

**List of Tables** 17

**List of Symbols** 19

**List of Acronyms & Abbreviations** 22

**1 Introduction** 26

1.1 The Electrification of the Transportation and Heating Sectors . . . 26

1.1.1 Electrifying Transport . . . 26

1.1.2 Electrifying Heat . . . 28

1.2 The Growing Share of Renewable-based Power Generation . . . 29

1.3 The Lack of Visibility in Distribution Networks . . . 31

1.4 The Need for Flexibility in Distribution Networks . . . 32

1.5 Power Electronics Devices as Enablers of Flexibility . . . 35

1.6 Research Aims . . . 36

1.7 Thesis Structure . . . 37

1.8 Contribution . . . 38

**2 Distribution Network Modelling** 40

2.1 Distribution Networks . . . 40

2.2 Power Flow . . . 44

2.3 Fixed-point Algorithm for Power Flow . . . 48
2.4 Optimal Power Flow . . . 52
2.4.1 Optimal Power Flow for Unbalanced Distribution Networks . . . 54
2.5 Summary . . . 57

**3 Active Network Management** **58**
3.1 Active Network Management . . . 58
3.2 Active Network Management Using Power Electronics Devices . . . 61
3.3 Active Network Management Control Using AC/DC OPF . . . 64
3.4 Case Studies . . . 68
3.4.1 IEEE 123-bus Test Feeder . . . 69
3.4.2 Lotherton and Canefields LV Substations . . . 80
3.4.3 European LV Network . . . 89
3.5 Limitations of the OPF for ANM control . . . 93
3.6 Alternatives to the OPF . . . 95
3.7 Summary . . . 97

**4 Data-driven Control for ANM** **99**
4.1 Regression . . . 99
4.2 Gradient-based Optimisation . . . 104
4.3 Case Studies . . . 105
4.3.1 IEEE 123-bus Test Feeder . . . 106
4.3.2 European LV Network . . . 110
4.4 Discussion . . . 113
4.5 Summary . . . 114

**5 Reinforcement Learning Control for ANM** **115**
5.1 Reinforcement Learning . . . 115
5.2 Model-free Reinforcement Learning for ANM . . . 117
5.2.1 Deep Q Learning Using Neural Networks . . . 120
5.2.2 Episodes and Termination Statements . . . 122
5.2.3 Case Study - Lotherton and Canefields LV Substations . . . 124
5.3 Model-based Reinforcement Learning for ANM . . . 131
5.3.1 Tools for Automatic Differentiation . . . 133
5.3.2 Adaptive Neural Networks for ANM . . . 136

5.3.3 Case Study - IEEE 123-bus Test Feeder - no Action Network . . . . . 138
5.3.4 Case Study - IEEE 123-bus Test Feeder - with Action Network . . . . . 143
5.4 Discussion . . . . . 146
5.5 Summary . . . . . 148

**6 Conclusions** **150**
6.1 Review of Problem Statement . . . . . 150
6.2 Contributions . . . . . 152
6.3 Conclusions . . . . . 154
6.4 Future Research . . . . . 158

**Bibliography** **160**

**Appendix A Fixed-point Algorithm for Power Flow** **175**
A.1 Lines and Cables . . . . . 175
A.2 Loads and Generators . . . . . 178
A.3 Fixed-point Algorithm . . . . . 179

# List of Figures


1.1 EVs global stock and selected countries [1]. . . . . 27

1.2 EV public charge points installed in the UK by DNO area (as of November 2020). Each dark blue marker denotes a charging point [14]. . . . . 28

1.3 Estimated heatpump uptake in the UK Power Networks area up to 2050, according to [17]. The bars represent the total number of heatpumps expected to be in operation under different scenarios. . . . . 29

1.4 Solar PV cumulative installations (top) and capacity in the UK (bottom) [19], from selected installation sizes only. The bars represent the total number of installations per each installation size range and the dashed lines illustrate their respective capacity. As of 2019 there were over 7GW of combined installed capacity from the four capacity ranges shown in the figure. . . . . 30

1.5 LCOE for utility-scale solar PV for selected countries and average [20]. . . . . 30

1.6 Illustration of the reduced visibility in the lower voltage levels of distribution networks. . . 31

1.7 Solar PV output over a week [22]. . . . . 32

1.8 Solar PV output over a month, all 31 days are superimposed [22]. . . . . 32

2.1 Typical power system structure with standard voltage levels for the UK. Large generation stations usually include several voltage transformation stages before connecting to the transmission network, for instance from 33 kV or 66 kV to 220 kV in the case of offshore windfarms [65]. . . . . 41

2.2 Distribution network structure with typical voltage levels for the UK [69]. The NOP can be used to connect two LV substations, via one of their feeders, to restore the supply in case of a fault. . . . . 42

2.3 Diagram of two LV feeders supplied from the same primary substation, with typical voltage levels for the UK. The usual single-phase connection of LCTs is illustrated. . . . . 44

2.4 Line connections between bus A and bus B and between bus B and bus C, depicting connections between internal nodes [69]. . . . . 49

3.1 Overview of network meshing in the OpenLV project. The monitoring units are shown in blue and the ALVIN Reclose™ devices used for ANM are in green [55]. . . . . 61

3.2 Two LV substations connected using power electronic converters, with the monitoring units shown in blue. This back-to-back connection between the converters is also referred to as SOP. . . . . 62

3.3 VSCs connected back-to-back through their respective DC bus forming a SOP. . . . . 65

3.4 Multiple VSCs forming a MTDC network. . . . . 67

3.5 IEEE 123-bus test feeder diagram [143, 144]. . . . . 69

3.6 Load and generation profiles with 1-hour resolution. . . . . 70

3.7 Voltage unbalance with default loads only. . . . . 70

3.8 Voltage unbalance with default loads, EVs and solar PV generators. . . . . 71

3.9 Current supplied by substation with default loads, EVs and solar PV generators. . . . . 71

3.10 Power supplied by substation with default loads, EVs and solar PV generators. . . . . 71

3.11 Active and reactive power losses by component with default loads, EVs and solar PV generators. . . . . 72

3.12 Runtime of the power flow for each time-step, the dashed line represents the average run-time. 72

3.13 IEEE 123-bus test feeder diagram with VSCs at marked buses (39, 66, 151, 300) [143, 144]. . . 73

3.14 IEEE 123 MTDC diagram. Each three-phase converter is composed of three single phase converters that share a common DC bus. . . . . 73

3.15 Voltage unbalance with default loads, EVs and solar PV generators with AC/DC OPF. . . . 74

3.16 Current supplied by substation with default loads, EVs and solar PV generators, calculated by the OPF. . . . . 74

3.17 Power supplied by substation with default loads, EVs and solar PV generators, calculated by the OPF. . . . . 75

3.18 Active power injected or absorbed by converters set out sequentially for the 24hr period. The red dashed line indicates the maximum active power that the converters can inject. . . 75

3.19 Reactive power injected or absorbed by converters set out sequentially for the 24hr period. The red dashed line indicates the maximum reactive power that the converters can inject. . 75

3.20 Active power exchanged internally, between the phases of each converter. . . . . 76

3.21 Active power exchanged between converters. . . . . 77

3.22 Active and reactive power losses by component with default loads, EVs and solar PV generators, as calculated by OPF. . . . . 77
3.23 Optimal power flow runtime for each time step, the dashed line represents the average runtime. The black solid circles represents the number of iterations required by each time step to converge. . . . . 78
3.24 Optimal power flow evolution of objective function for each time step. . . . . 78
3.25 Optimal power flow evolution of objective function for each time step. . . . . 79
3.26 Optimal power flow evolution of the converters' power injections for hour 7. The power injections are set sequentially following the number of iterations required for the OPF to converge. . . . . 79
3.27 Optimal power flow evolution of the converters' power injections for hour 14. The power injections are set sequentially following the number of iterations required. . . . . 79
3.28 Overview of network meshing in the OpenLV project. The monitoring units are shown in blue and the ALVIN Reclose™ devices used for ANM are in green [55] (repeated here for clarity). . . . . 80
3.29 Network map image of Lotherton and Canefields substations. The square on the upper right corner and on the bottom left corner show where the LV substations are located. The blue lines depict the LV and service cables used to supply the customer properties. Reproduced with permission from Western Power Distribution [148]. . . . . 81
3.30 Aerial image of Lotherton and Canefields substations. The yellow line illustrates the path followed by the underground feeders of each substation and the red square illustrates the location of the NOP. . . . . 82
3.31 Aerial image of Lotherton and Canefields substations with the network map overlaid. The blue lines depict the LV and service cables used to supply the customer properties and the red square illustrates the location of the NOP. Reproduced with permission from Western Power Distribution [148]. . . . . 82
3.32 Total active power supplied by each substation, Lotherton (blue) and Canefields (red), over a sample week in November. . . . . 83
3.33 Current supplied by Lotherton substation for November. The solid line depicts the average for that month and the shaded area represents the standard deviation. . . . . 83
3.34 Current supplied by Canefields substation for November. The solid line depicts the average for that month and the shaded area represents the standard deviation. . . . . 84
3.35 Daily samples and monthly average of total power flowing through the ALVIN Reclose™ units. . . . . 84

3.36 Actuation times of the ALVIN Reclose™ units are indicated by the solid line. . . . . 85
3.37 Two LV substations connected using power converters, with the monitoring units shown in blue. This back-to-back connection between the converters is also referred to as SOP (figure repeated here for clarity). . . . . 86
3.38 Comparison of the transformers loading and their difference. The solid lines indicate the results of the base case, whereas the dashed and dotted lines indicate the results of the switching and OPF, respectively. The bars indicate the loading difference between the transformers. 86
3.39 Comparison of the current flowing through Feeder 3 and Feeder 2 for the base case, with the ALVIN Reclose™ units and with the converters. . . . . 87
3.40 Comparison of transformers loading for each case. . . . . 88
3.41 Comparison of daily energy losses by network component. From left to right, the bars indicate the case without ANM first, followed by the Alvin case and the OPF. . . . . 88
3.42 Diagram of the European low voltage network. The substation is highlighted by a black triangle, whereas the loads are depicted by the blue circles. Those customers with EVs have a yellow circle overlayed and those with solar PV have a diamond marker. The potential location of the VSCs is illustrated by the red squares [143, 144]. . . . . 90
3.43 Run-time of the power flow for each time-step, the dashed line represents the average runtime. 91
3.44 Current supplied by substation with default loads, EVs and solar PV generators. . . . . 91
3.45 Optimal power flow runtime for each time step, the dashed line represents the average runtime. The black solid circles represents the number of iterations required by each time step to converge. . . . . 92
3.46 Current supplied by substation with default loads, EVs and solar PV generators, calculated by the OPF. . . . . 92
3.47 Optimal power flow evolution of objective function for each time step. . . . . 93
3.48 Active power injected or absorbed by converters set out sequentially for the 24hr period. The red dashed line indicates the maximum active power that the converters can inject. . . 93
3.49 Overview of the typical monitoring available in LV networks. The grey shaded area indicates the lack of visibility of individual buses and loads along the outgoing LV feeders. . . . 95

4.1 Two LV substations connected using power electronic converters, with the monitoring units shown in blue. This back-to-back connection between the converters is also referred to as SOP. . . . . 100
4.2 Illustration of the range of values produced by the tanh activation function. . . . . 103

4.3 Neural network with 4 inputs, 2 outputs and a single hidden layer with 8 hidden units. This neural network has 48 weights [170]. . . . . 103
4.4 Overview of the regression training and use of the data-driven control method. . . . . 105
4.5 Load and generation profiles with 20-minute resolution. . . . . 105
4.6 IEEE 123-bus test feeder diagram with VSCs at marked buses (39, 66, 151, 300) [143, 144]. . . 106
4.7 Histogram of error of neural network regression model with respect to the OPF for the IEEE 123-bus test feeder. . . . . 108
4.8 Comparison of active power set-points from regression model and OPF of each converter phase (set out sequentially for the 24hr period). . . . . 108
4.9 Comparison of the current at head of feeder and voltage unbalance factor for the IEEE 123-bus test feeder. . . . . 109
4.10 European LV network diagram with VSCs at marked buses [175]. . . . . 110
4.11 Comparison of active power set-points from regression model and OPF of each converter phase (set out sequentially for the 24hr period). . . . . 111
4.12 Comparison of active power set-points from regression model and OPF of each converter phase (set out sequentially for the 24hr period). . . . . 112
4.13 Current at head of feeder 2 for base case, OPF and regression model. . . . . 113

5.1 Agent-environment interaction over multiple time-steps in a Markov process [181]. . . . . 117
5.2 Overview of the model-free training environment for an ANM. The substation transformer and outgoing feeders are modelled in OpenDSS. The aggregated load supplied by the feeders can be obtained from monitoring devices at the LV substation (e.g., monitoring devices installed in the OpenLV project [195]). The interactions between the agent and the environment are performed via power flow simulations and are used to train the neural network to make optimal decisions. . . . . 118
5.3 Overview of the proposed control based on deep Q learning. The LV network, composed of two substations and monitoring devices at the LV side of the transformer and at the head of the outgoing feeders, represent the environment. The power converters represent the learning agent and their power injections are the actions. The neural network approximates the value of the state-action pairs. . . . . 121
5.4 Evolution of the power injections of each phase of the two power converters and the loading of the two substation transformers which are to be equalised. These are the iterations at time-step 200. The dotted and dashed lines represent the converters' power injections, the solid lines indicate the transformers loading. . . . . 125

5.5 Evolution of the cumulative reward and the transformers' power losses for time-step 200. . 126
5.6 Q values of actions averaged over total iterations of each of the 288 time-steps. . . . . 126
5.7 Comparison of active power set-points from the OPF and the DQN for each converter phase for Sunday (set out sequentially for the 24hr period). . . . . 127
5.8 Histogram of the error between the injections calculated by the DQN and those by the OPF for the base case. The mean of the errors, the RMSE, and the MAE are included as annotations. . . . . 127
5.9 Comparison of transformer loading for the base case, equalisation using an OPF and the actions of the DQN control agent. . . . . 128
5.10 Total power losses comparison between base case, OPF, and DQN. The bars represent the power losses at each time-step for the base case; the solid and dotted lines represent the power losses for the DQN and OPF, respectively. . . . . 128
5.11 Load profiles at the head of the feeders and their means for Sunday. . . . . 129
5.12 Load profiles at the head of the feeders and their means for Monday. . . . . 129
5.13 Comparison of active power set-points from the OPF and the DQN for each converter phase for Monday (set out sequentially for the 24hr period). . . . . 130
5.14 Histogram of the error between the injections calculated by the DQN and those by the OPF for the Monday load profile. The mean of the errors, the RMSE and the MAE are included as annotations. . . . . 130
5.15 Transformers' loading comparison between the base case, DQN, and the OPF for the Monday load profile. . . . . 131
5.16 Example of a simple computational graph with two variables and one constant [210]. . . . . 134
5.17 Example of a computational graph in TensorFlow. . . . . 135
5.18 Computational graph with intermediate results to illustrate the reverse mode of AD [211]. . 135
5.19 Overview of the model-based RL optimisation environment. The network and the load profiles are modelled using the power flow equations, which solved via the fixed-point iteration method. The actions are optimised directly by gradient descent. . . . . 137
5.20 Overview of the model-based optimisation environment, similar to Fig. 5.19 but including the action network, which produces the actions. . . . . 138
5.21 Total daily energy losses at every training iteration. The baseline is shown by the red dashed line whereas the minimum reached by the OPF is shown by the blue dotted line. . . . . 139
5.22 Active and reactive power losses by component type. . . . . 139

5.23 Value of the gradients during the first and last training iteration. Each of the colours represent the gradient of the total power losses with respect to one of the actions (i.e., the converters power injections). . . . . 139
5.24 Active power injected or absorbed by converters set out sequentially for the 24hr period. The different colour of the bars indicate the phases. The red dashed line indicates the maximum active power that the converters can inject. . . . . 140
5.25 Reactive power injected or absorbed by converters set out sequentially for the 24hr period. The different colour of the bars indicate the phases. The red dashed line indicates the maximum active power that the converters can inject. . . . . 140
5.26 Comparison of the current supplied by the substation, for the baseline and for the optimised case. . . . . 141
5.27 Comparison of active power set-points from the OPF and the Gradient Descent optimisation for each converter phase (set out sequentially for the 24hr period). . . . . 141
5.28 Histogram of the error between the injections calculated by the DQN and those by the OPF for the Monday load profile. The mean of the errors, the RMSE and the MAE are included as annotations. . . . . 142
5.29 Runtime for each operation. For the training iterations, the red bar indicates the time taken by the first training iteration, the blue bar indicates the time taken by the subsequent 19 iterations. . . . . 142
5.30 Illustration of the neural network architecture, comprising 18 inputs (at the top), 36 outputs (bottom) and 2 hidden layers [170]. This neural network has more than 9,000 trainable weights. . . . . 143
5.31 Total energy losses at every training iteration. The baseline is shown by the red dashed line whereas the minimum reached by the OPF is shown by the blue dotted line. . . . . 144
5.32 Active power injected or absorbed by converters set out sequentially for the 24hr period. The different colour of the bars indicate the phases. The red dashed line indicates the maximum active power that the converters can inject. The optimal power injections are obtained by the neural network (actor). . . . . 144
5.33 Reactive power injected or absorbed by converters set out sequentially for the 24hr period. The different colour of the bars indicate the phases. The red dashed line indicates the maximum active power that the converters can inject. The optimal power injections are obtained by the neural network (actor). . . . . 144
5.34 Comparison of the current supplied by the substation, for the baseline and for the optimised case. . . . . 145

5.35 Runtime for each operation. For the training iterations, the red bar indicates the time taken by the first training iteration, the blue bar indicates the time taken by the subsequent 24 iterations. . . . . 145

# List of Tables


1.1 Estimated cost for replacing cables in distribution networks [36, 37]. . . . . 33

2.1 Typical LV cables data. . . . . 42

2.2 Typical parameters for transformers used in LV networks [69, 79]. . . . . 43

3.1 Network data and peak load information. . . . . 85

3.2 Cases with increased LCTs penetration. . . . . 88

3.3 Details of the modified European LV network. . . . . 89

3.4 Base case results European LV network for the daily load profile. . . . . 91

3.5 Results comparison European LV network. . . . . 92

4.1 Network measurements used for training the regression model. . . . . 107

4.2 Training and validation results for the four regression methods. . . . . 107

4.3 Comparison of regression method performance against OPF and base case for IEEE 123-bus test feeder. . . . . 109

4.4 Training and validation results for the EULV network. . . . . 111

4.5 Comparison of regression model performance in EULV Network. . . . . 112

4.6 Comparison of regression model performance in EULV Network - Feeder 1 and 2. . . . . 112

5.1 Summary of LV network measurements. . . . . 121

5.2 Discretized action space sample. . . . . 122

5.3 Summary of the parameters used for training. . . . . 125

5.4 Results comparison for entire day - Training. . . . . 129

5.5 Results comparison for entire day - Validation. . . . . 131

5.6 Results and runtime per-time step IEEE 123-bus test feeder. . . . . 146

6.1 Comparison between alternative methods proposed and traditional optimisation. . . . . 157

A.1 Overhead line configurations. . . . . 175

A.2 Underground cables configuration. . . . . 175

# List of Symbols

| Notation | Description |
|---|---|
| $a_s$ | Phasor rotation operator |
| $\alpha$ | Learning rate |
| $B_{ij}$ | Susceptance of a line between nodes $i$ and $j$ |
| $\hat{\beta}_0$ | Intercept of the estimated regression model |
| $\hat{\beta}_j$ | Coefficients estimated by the regression model |
| $\Delta$ | Connection in delta |
| $D_{ij}$ | Spacing between conductors $i$ and $j$ |
| $E_i$ | Complex voltage phasor at a node $i$ in rectangular form |
| $\mathbb{E}$ | Denotes expectation of a random variable |
| $\epsilon_t$ | Exploration-exploitation ratio |
| $f$ | Frequency |
| $G_{ij}$ | Conductance of a line between nodes $i$ and $j$ |
| $\gamma$ | Discount-rate parameter |
| $I_{comp,element}$ | Compensation current of an element |
| $I_{linear,element}$ | Linear part of the current of an element |
| $I_{VSC,i}$ | Internal current of the converter $i$ |
| $\Delta I_i$ | Current mismatch for node $i$ |
| $\Delta I_{Re,i}$ | Real part of the current mismatch for node $i$ |
| $\Delta I_{Im,i}$ | Imaginary part of the current mismatch for node $i$ |
| $k$ | Number of iterations |
| $\nabla_x$ | Gradient of a function with respect to x |
| $\Omega$ | ohm (unit) |
| $\pi$ | A policy to be followed by a learning agent |

| Notation | Description |
|---|---|
| $P_i$ | Active power injection at node $i$ |
| $P_{g,i}$ | Active power injection of generator at node $i$ |
| $P_{d,i}$ | Active power injection of load at node $i$ |
| $P_i^{sp}$ | Specified active power injection at node $i$ |
| $P_{injVSC,i}$ | Active power injected into the AC side of the converter $i$ |
| $P_{lossVSC,i}$ | Active power losses of the converter $i$ |
| $P_{DC,i}$ | Active power injected into the DC side of converter $i$ |
| $P_{l,m}$ | Power limit at line $m$ |
| $Q(s,a)$ | Action-value function for action $a$ in state $s$ |
| $Q_i$ | Reactive power injection at node $i$ |
| $Q_{injVSC,i}$ | Reactive power injected into the AC side of the converter $i$ |
| $Q_{g,i}$ | Reactive power injection of generator at node $i$ |
| $Q_{d,i}$ | Reactive power injection of load at node $i$ |
| $Q_i^{sp}$ | Specified reactive power injection at node $i$ |
| $R^2$ | coefficient of determination |
| $\rho$ | earth resistivity |
| $r_i$ | resistance of conductor $i$ |
| $S_{base}$ | Apparent power per-unit base |
| $S_i$ | Apparent power injection at node $i$ |
| $\tau$ | Exploration decay ratio |
| $\theta_i$ | Voltage angle at node $i$ |
| $\theta$ | Weights of a neural network |
| $V_1$ | Positive sequence component of voltage |
| $V_2$ | Negative sequence component of voltage |
| $V_{AC,i}$ | voltage at the AC bus of converter $i$ |
| $V_i^k$ | Voltage at node $i$ iteration $k$ |
| $V_{base}$ | Voltage per-unit base |
| $V_{DC,i}$ | Voltage at the DC bus of converter $i$ |
| $V_{error}$ | Error between the previous and the present voltage iteration |
| $V_{Im,i}$ | Imaginary part of the complex voltage phasor at node $i$ |

| Notation | Description |
|---|---|
| $V_{Re,i}$ | Real part of the complex voltage phasor at node $i$ |
| $Z_{base}$ | Impedance per-unit base |
| $Y_{ij}$ | Nodal admittance between nodes $i$ and $j$ |
| $Y_{prim}$ | Primitive admittance matrix |
| $Y_{system}$ | System admittance matrix |
| $Z_0$ | Zero sequence impedance |
| $Z_1$ | Positive sequence impedance |
| $Z_2$ | Negative sequence impedance |
| $Z_{ii}$ | Primitive self-impedance of conductor $i$ |
| $Z_{ij}$ | Primitive mutual-impedance between conductors $i$ and $j$ |
| $Z_{in}$ | Primitive mutual-impedance between conductor $i$ and the neutral $n$ |
| $Z_{prim}$ | primitive impedance of a line |
| $Z_{abc}$ | phase impedance of a line |
| $Z_{012}$ | sequence impedance of a line |

# List of Acronyms & Abbreviations

**A** Ampere

**AC** Alternating Current

**ACSR** Aluminium conductor steel-reinforced

**AD** Automatic Differentiation

**Adam** Adaptive Moment Estimation

**ADMD** after diversity maximum demand

**ADMM** Alternating Direction Method of Multipliers

**ADP** Adaptive Dynamic Programming

**AGC** automatic generation control

**ANM** Active Network Management

**BPTT** Back-propagation Through Time

**BSP** Bulk Supply Point

**CART** Classification and Regression Trees

**CIGRE** *Conseil International des Grands Réseaux Électriques*

**CIM** Current Injection Method

**DC** Direct Current

**DER** Distributed Energy Resources

**DG** Distributed Generation

**DNO** Distribution Network Operator

**DQN** Deep Q Network

**DSO** Distribution System Operator

**DSR** Demand Side Response

**DUoS** Distribution Use of System

**EHV** Extra High-voltage

**ENA** Energy Networks Association

**EULV** European Low Voltage

**EV** Electric Vehicle

**EVCP** Electric Vehicle Charge Point

**FACTS** Flexible Alternating Current Transmission System

**FiT** Feed-in Tariff

**GBT** Gradient boosted trees

**GMR** Geometric Mean Radius

**GSP** Grid Supply Point

**HV** High-voltage

**HVDC** High-voltage Direct Current

**Hz** Hertz

**IEEE** Institute of Electrical and Electronic Engineers

**IPM** Interior Point Method

**km** Kilometer

**kV** Kilovolts

**kVA** Kilovolt Ampere

**kvar** Kilovolt Ampere Reactive

**kW** Kilowatt

**kWh** Kilowatt-hour

**LCOE** Levelised Cost of Electricity

**LCT** Low-Carbon Technology

**LV** Low-Voltage

**MV** Medium-Voltage

**MVA** Megavolt Ampere

**MW** Megawatt

**MDP** Markov Decision Process

**MLP** Multi-layer Perceptron

**MSE** Mean Square Error

**MTDC** multi-terminal DC

**NOP** Normally-Open Point

**OPF** Optimal Power Flow

**P2P** Peer-to-peer

**PCC** Point of Common Coupling

**PED** Power Electronics Device

**POC** Point of Connection

**PV** Photovoltaic

**RF** Random Forests

**RL** Reinforcement Learning

**RMS** Root Mean Square

**RSS** Residual Sum of Squares

**RTU** Remote Terminal Unit

**SCADA** Supervisory Control and Data Acquisition

**SDP** Semi-definite programming

**SOCP** second order cone programming

**SOP** Soft-Open-Point

**TOU** time-of-use

**UK** United Kingdom

**USD** United States Dollar

**V** Volts

**VA** Volt Ampere

**var** Volt Ampere Reactive

**VUF** Voltage Unbalance Factor

**W** Watt

**VSC** Voltage Source Converter

# Chapter 1

# Introduction

Distribution network operators (DNOs) face several challenges over the coming years: a) the electrification of both the transportation and the heating sectors, b) the growing number of renewable-based distributed generation systems being installed in low- and medium-voltage networks and c) the low visibility of the lower voltage levels of the network, due to the lack of instrumentation and monitoring. This chapter provides evidence of the changing landscape in distribution networks, the challenges associated with the integration of low-carbon technologies and the impact these technologies have on the planning and operation of these networks. It then draws attention to the urgent need for *flexibility* required by the DNOs, which can come from the customer demand and generation assets but can also be provided by active network management schemes, enabling a more comprehensive controllability of the power flows in these networks.

## 1.1 The Electrification of the Transportation and Heating Sectors

### 1.1.1 Electrifying Transport

The past few years have seen a growing trend towards the electrification of the transportation sector. According to the IEA, in less than ten years the global electric vehicle (EV) stock increased from a few hundred thousands to more than 7 million [1], as can be observed in Fig 1.1. This was supported in part by low-carbon policies and incentives, and also by the continuous decline in battery costs [1,2]. A recent study [3] estimated that, for car manufacturers, the current (as of 2019) price of battery packs is between 210-230 USD/kWh and could potentially reach 150 USD/kWh by the end of 2020. This price point would make the production cost of EVs with a 200-300 mile range comparable to that of conventional vehicles.

Considering this trend, it is clear that the number of EVs will grow more rapidly than previously anticipated [4], especially in urban centres. Shifting this large energy demand from liquid fuels to electricity will bring challenges to power systems. The magnitude of the power demand for even a small number of vehi-

cles can become significant and small clusters of EVs can be difficult for distribution networks to manage. EVs represent one of the largest DC-based loads in low-voltage (LV) feeders, where the after diversity maximum demand (ADMD) for residential customers is around 2kW (assuming gas is used for heating) [5, 6]. The ADMD is calculated using measurements for a year by aggregating the demand over a large number of customers and dividing it by the total number of customers, thus it takes into account the coincident peak load [5–7]. Level 1 chargers for EVs, mostly used at residential and workplace locations in the United Kingdom, can absorb up to 2.3 kW, whilst a DC fast charger, usually located at a dedicated charging station, can draw up to 130 kW [8, 9]. This means that 1,000 EVs charging simultaneously in an urban area would increase the power demand by 2.3 MW, assuming the lowest charging level (2.3 kW) is used. If EV adoption rates continue to accelerate, the additional demand coming from these vehicles could represent 8% of the total electricity demand in the next two decades [10].

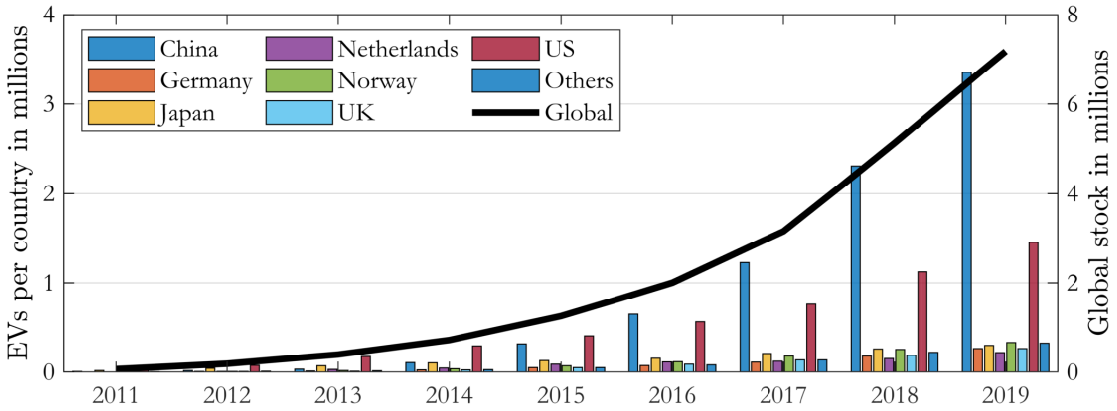



Figure 1.1: EVs global stock and selected countries [1].

On top of this, the study [11] pointed out that automated or autonomous vehicles could increase vehicle utilisation rates, as they could make trips without requiring a driver, thus making EVs even more cost-effective over their lifetime. Whilst this could eventually translate into fewer vehicles (e.g. vehicles could be shared), it would increase the rate at which they charge/discharge, putting additional strain on electricity networks. A more recent report [12] predicts that with the transition from the current ownership model to a Transport-as-a-Service (TaaS) model, these autonomous EVs will be used more than twice as much as non-EVs, some of them reaching over 500,000 miles driven, compared to an average of 140,000 miles for current vehicles. Moreover, the increased availability of fast charging stations will lead to heightened demand peaks throughout the day [13]. The total number of public EV charging points (EVCPs) installed in the UK (14,198 as of November 2020), according to data from Open Charge Map [14], is shown in Fig 1.2. It can be seen that the large majority are located in the Greater London area, as well as in major cities in the UK.

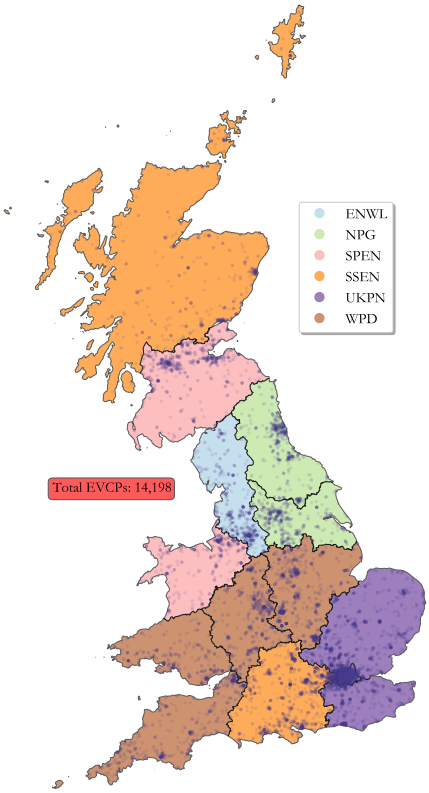



Figure 1.2: EV public charge points installed in the UK by DNO area (as of November 2020). Each dark blue marker denotes a charging point [14].

### 1.1.2 Electrifying Heat

The report in [15] estimated that water and space heating represent approximately 23% of the total energy demand in the UK and are responsible for around 20% of the total carbon emissions, approximately 100 $MtCO_2$ annually. Hence, to fully decarbonise the energy system, the heating sector also needs to transition to more efficient and low-carbon alternatives than gas, such as electricity or a mix of electricity and hydrogen.

Electrifying the heat demand can be done through two main technologies, heat pumps and direct electric resistive heaters. Whilst direct heating has a lower initial cost for customers, it is less efficient than heat pumps, resulting in a higher demand [16]. Nonetheless, the requirement for thermally efficient buildings needs to be taken into account in the case of heat pumps. Although, deployment of heating has lagged behind other technologies, such as EVs, it is expected that the uptake will accelerate in the next two decades. The report in [17] produced several scenarios estimating the number of heatpumps in the UK Power Networks area, ranging from 4 to 7 million, as illustrated in Fig 1.3. This will translate into an additional demand pressure for distribution networks.

A recent report analysing the infrastructure cost of decarbonising heat in the UK estimated that the electrification of heat would translate into an additional 45 GW of demand and would require a major investment (of approximately £20bn) in both distribution and transmission networks over the next 30 years [16] A

more detailed analysis of the heating demand was carried out in [18]. Where the authors developed a model to estimate the half-hourly demand for domestic heating during mild, normal and cold years. The annual peak demand for the years modelled was estimated at 129 GW, 131 GW and 141 GW, for mild, normal and cold respectively.

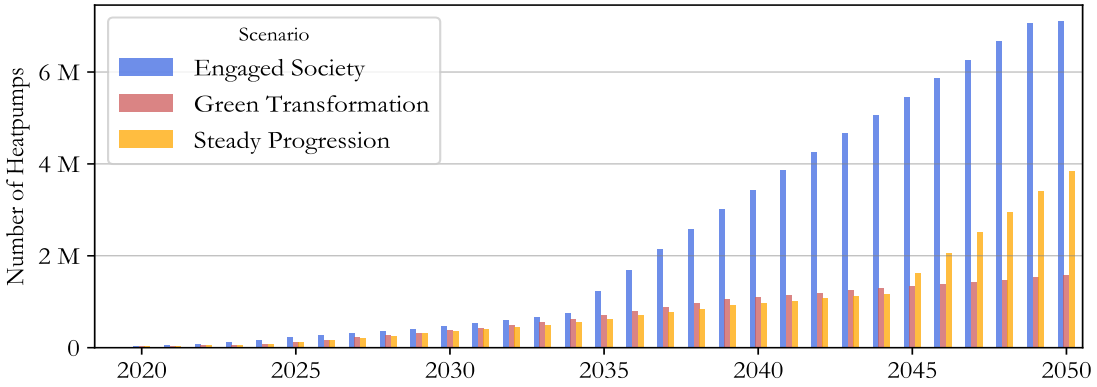



Figure 1.3: Estimated heatpump uptake in the UK Power Networks area up to 2050, according to [17]. The bars represent the total number of heatpumps expected to be in operation under different scenarios.

The increased peak demand coming from the electrification of both the transportation and the heating sector can be challenging for DNOs. In particular, if the peak demand from these technologies coincides with each other and sits on top of the usual evening peak, DNOs might need to upgrade transformers and cables in certain parts of the network.

## 1.2 The Growing Share of Renewable-based Power Generation

Electricity networks are experiencing a transition, at least in part, from a centralised model of power generation, transmission, and distribution, towards a decentralised, consumer-driven system based on renewable resources. The installed capacity of solar PV has grown significantly, as can be observed from Fig. 1.4 for the UK [19], underpinned by the efforts from countries to move away from fossil fuels and by the dramatic fall in costs observed since 2010. The global levelised cost of electricity (LCOE) from utility-scale solar PV has declined approximately 77% between 2010 and 2018, as shown in Fig. 1.5 [20].

Policy support, in the form of incentives, such as the Feed-in-Tariff (FiT) in the UK, has underpinned the rise of the consumer/producer, usually called *prosumers*. Fig. 1.4 shows that, from the total installed capacity of solar PV, a significant proportion comes from small-scale installations between 0-4 kW, with approximately 1 million installations comprising almost half of the total capacity. As of 2019, these small-scale solar PV installations totalled 2.6 GW of capacity, whereas 3.5GW was generated by over 4,000 installations of between 0.05-5 MW. The FiT also promoted the growth of "community energy" projects [21]. These

projects are usually carried out by social enterprises and small organisations, with the goal of democratising the access to renewable energy and lowering consumer energy costs.

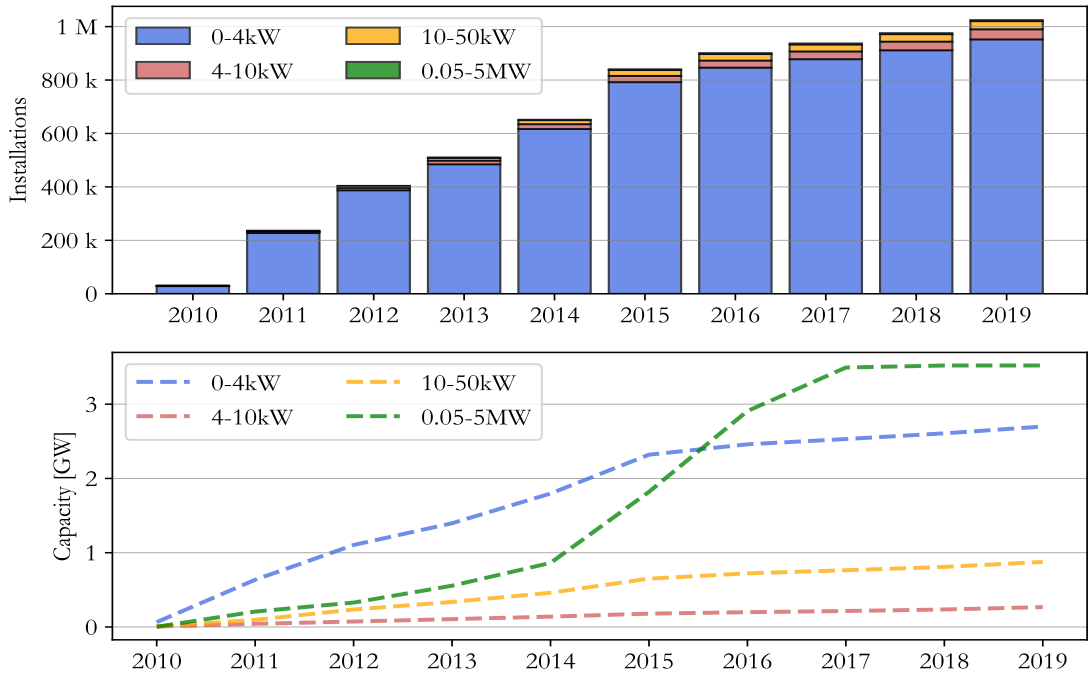



Figure 1.4: Solar PV cumulative installations (top) and capacity in the UK (bottom) [19], from selected installation sizes only. The bars represent the total number of installations per each installation size range and the dashed lines illustrate their respective capacity. As of 2019 there were over 7GW of combined installed capacity from the four capacity ranges shown in the figure.

The authors of [21] surveyed such projects and reported that the majority were rooftop solar PV, with an average generating capacity of 74 kW. Whilst these projects are often scattered over many buildings, the overall capacity is significant compared to the typical small-scale residential installations of several kW [22]. Despite the incentives being slowly scaled back, the continuous decline in installation cost for solar PV and its combination with other technologies, such as small-scale energy storage, might be sufficient to keep some of these projects viable.

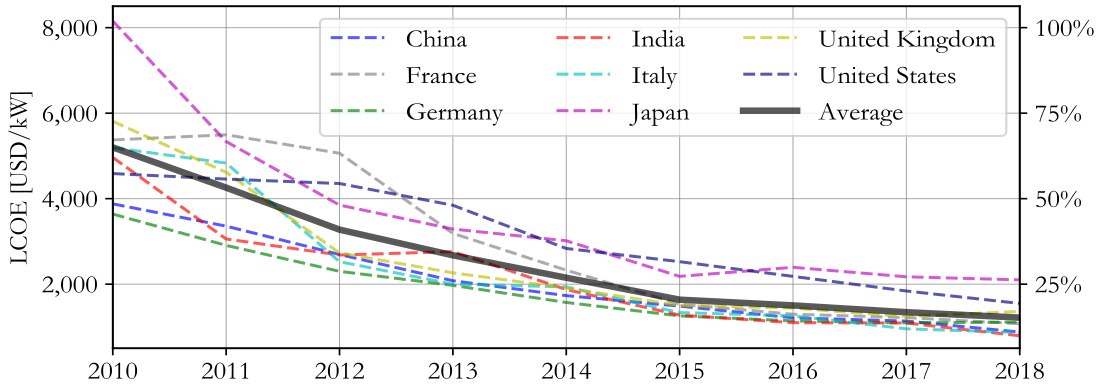



Figure 1.5: LCOE for utility-scale solar PV for selected countries and average [20].

One of the main challenges associated with solar PV generation is that its output depends on the size of the installation, the weather and time of day. This can reduce the visibility of actual power demand in distribution networks, as the amount of energy needed by the customer is uncertain [23]. This requires significant power changes to meet the power demand at night when there is no output from solar PV units [24], which might be worsened by the connection of EVs and heatpumps. Unless smart meters are capable of discerning between the load and generation, or forecasting algorithms are used, this presents a problem for DNOs in terms of network visibility.

## 1.3 The Lack of Visibility in Distribution Networks

Historically, the coverage by monitoring and instrumentation devices in distribution networks, and in particular LV networks, has been low. One of the reasons is the segmentation and extensive nature of these networks, which makes it prohibitively expensive to monitor every node [25]. Also, monitoring LV networks was not required in the past, given the predictability of the loads and their growth [26]. Moreover, the power flows were unidirectional due to the lack of embedded generation at this voltage level. Apart from the growing number of smart meters being installed at customers locations [27], the observability of LV substations and feeders remains low compared to the transmission system, as illustrated in Fig. 1.6. Load profiles representative of the different seasons are normally used to forecast the expected load [28, 29].

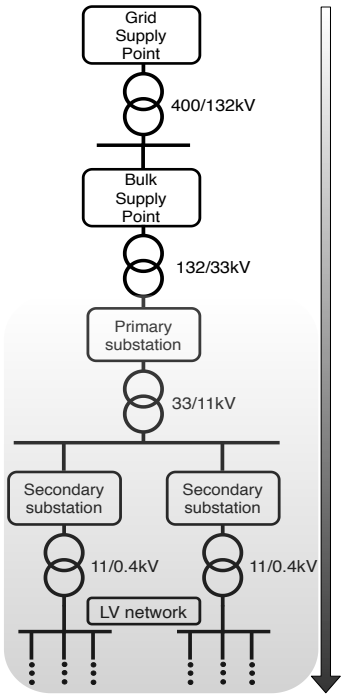



Figure 1.6: Illustration of the reduced visibility in the lower voltage levels of distribution networks.

Moreover, in many cases, the actual topology and connectivity of the network is only partially known. The phase-connectivity of the customers within the LV network is not usually specified, making it harder

to predict which phase will be the most affected and making the impact of LCTs harder to assess in detail. Increasing the visibility of LV networks could have a key role in supporting targeted deployment of smart grid solutions to reduce or defer the cost of network reinforcement.

## 1.4 The Need for Flexibility in Distribution Networks

Distribution networks will be the part of the electricity system which will present the greatest number of problems in the medium and long term, because of the increasing amount of low-carbon technologies (LCTs) connected at these voltage levels and the lack of active control of the network [30]. The large power variations will result in several operational problems, that will impact the reliability of these networks. Rapid voltage fluctuations, increased thermal constraints, greater phase imbalance, and increased energy losses are among the main concerns for DNOs [31, 32]. Fig. 1.7 and Fig. 1.8 depict the power output of a solar installation in London, over a week and over a month, respectively [22].

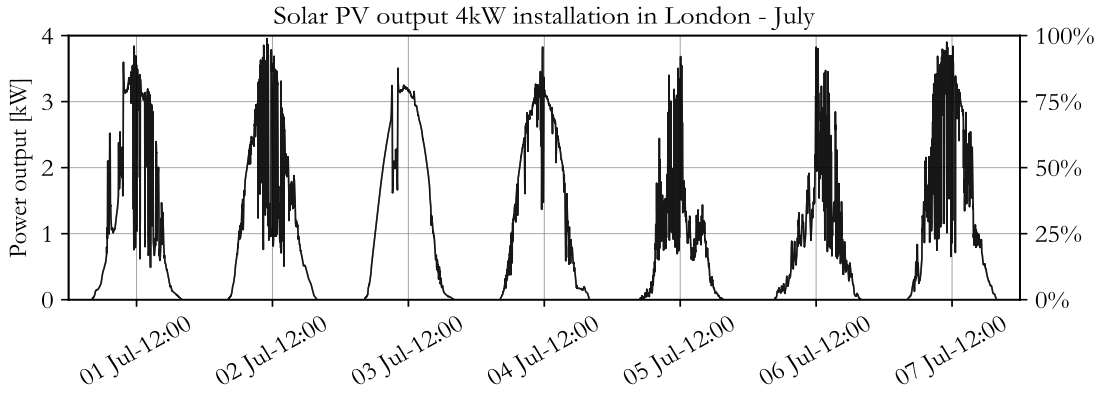



Figure 1.7: Solar PV output over a week [22].

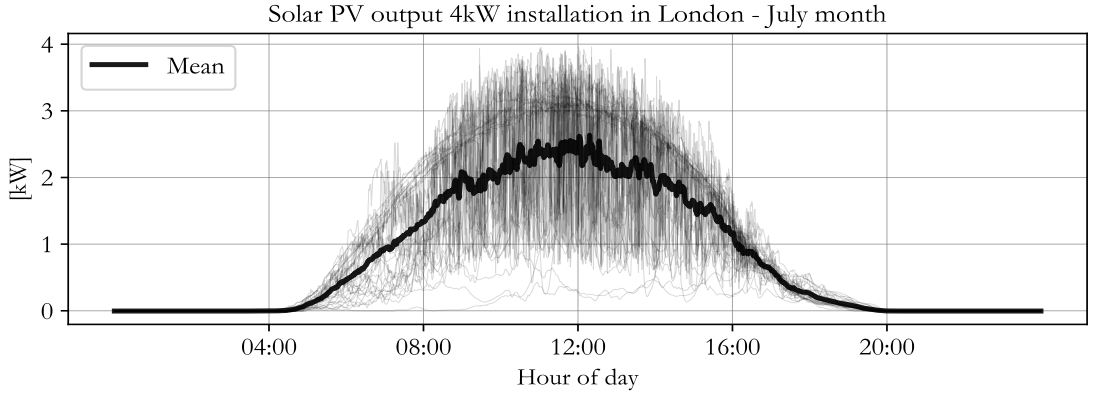



Figure 1.8: Solar PV output over a month, all 31 days are superimposed [22].

These figures illustrate the sudden changes in power output that occur in a solar PV installation due to weather variations over the course of the day. From Fig. 1.7 it can be observed that, within minutes, the

power output can change from full output (100%) to less than 25% of the rated output, due to a change in weather conditions; whereas Fig. 1.8 shows that, although the monthly average behaviour between days is broadly similar, the daily behaviour is erratic, marked by peaks and troughs.

As a consequence, DNOs will need to find alternative solutions to avoid or defer large-scale network reinforcement. Recently, one of the DNOs in the UK reported having over 8.5 GW of distributed generation connected to its network and outlined some of the difficulties in managing the network constraints [33].

Traditionally, reinforcement and upgrades of distribution networks are done based on load growth forecasts or after incidents occur, such as fuse operations, exceeded voltage levels, or customer disconnections [34, 35]. However, with the increasing volume of LCTs, the number of such incidents can rise significantly, making this approach costly and thus less appealing. Table 1.1 shows the estimated cost per km for replacing cables in distribution networks according to the Office of Gas and Electricity Markets (Ofgem) [36, 37].

Table 1.1: Estimated cost for replacing cables in distribution networks [36, 37].

| **Cable type** | **Cost/km** |
|---|---|
| LV main (Underground plastic) | £98.4k |
| 6.6/11 kV Underground Cable | £82.9k |
| 33 kV Underground Cable | £256.8k |
| 132 kV Underground Cable | £1,047.1k |

Although investment in network infrastructure is normally made by the DNOs, this cost is then transferred to the energy suppliers through Distribution Use of System (DUoS) tariffs. Ultimately, these costs are charged to consumers, as a component of their energy bill. Thus, enabling the connection of more LCTs whilst keeping network reinforcements cost-effective is in the interest of all parties, from generators to DNOs and consumers [38].

An additional issue to consider is that three-phase line currents and voltages in distribution networks are often unbalanced because of the single-phase connection of loads and generators [39]. Given that this unbalance varies dynamically throughout the day (i.e. the demand from each load varies independently from one another), LCTs would only make the situation worse for already unbalanced networks. A large voltage unbalance in the network can reduce the life of the transformer at the substation due to the concentration of loads in one of the phases and its impact on the insulation [40]. Because the temperature distribution would no longer be uniform, the insulation of the phase operating at the highest temperature would suffer greater deterioration.

In the past, the methods used by many utilities or DNOs to screen new requests for the connection of distributed generation (DG), such as solar PV, relied on the relation between the amount of power to be connected and the peak load, also referred to as capacity penetration [41]. However, several recent reports have raised concerns regarding the limitations of such screening methods and have highlighted the need for

more tailored approaches that take into account the individual characteristics of the distribution feeders [42]. The term *hosting capacity* was introduced to fill this gap. The Electric Power Research Institute (EPRI) defines *hosting capacity* as the amount of distributed energy resources (DERs) that can be incorporated into a network under existing control and infrastructure configurations, without having an adverse impact on the power quality, the reliability, or the ageing of the distribution equipment [41, 43]. For instance, over-voltage/under-voltage events in the primary or secondary systems, malfunction of the protection system, or thermal overload of network components. In the case of distributed solar PV, the National Renewable Energy Laboratory (NREL) states that the hosting capacity can be seen as the amount of peak output power (in kW or MW) that can be installed along a certain feeder without resulting in a breach of an operating limit [41]. This hosting capacity varies from one feeder to another, depending on characteristics such as the voltage level, the short-circuit fault current, the $X : R$ ratio, the presence of voltage regulators, the loading level, and the location of the DER resource within the feeder [41]. Other studies define hosting capacity in terms of the aggregated penetration level of a LCT considering both LV and MV networks [44]. Thus the hosting capacity is measured as the the percentage of customers within a network that can install solar PV without overloading the substation transformer or incurring in over-voltages at the customer nodes.

As well as generators, EVs can also present challenges to network operators. A recent study by a utility company in California [45] revealed that uncontrolled EV charging could increase the peak demand in residential feeders by 64% on average; this peak had resulted in approximately 60% of the feeders requiring some form of reinforcement. The authors assumed that each house would have one EV and that level 1 and level 2 charging was used either at home or at the workplace. The study conclusions highlighted the need for managing the EV charging process via time-of-use (TOU) tariffs. Another alternative would be the implementation of demand side response (DSR) schemes, such as the one proposed in [46], where the customers can participate by shifting their load following a request from the DNO. A study focusing on typical networks in the UK found that voltage would exceed the lower limit during the winter peak and thermal problems would occur in both the LV feeders and the substation transformer for medium and high EVs uptake scenarios [47]. Another recent study [32] analysed the impact of EV charging on the life of distribution transformers in residential areas. This study employed high-resolution data from more than 1,000 distribution transformers with different loading characteristics, as well as load profiles from residential consumers and EV charging data. The results showed a large increase in overloaded (1.25-1.5pu) and critically overloaded (1.5pu) transformers for penetration levels of more than 10%. The exact topology of the network was not taken into account by the study, hence, the actual results might be worse than those modelled. Furthermore, it should be noted that the majority of EVs in distribution networks are connected to a single phase, adding to the network phase imbalance problem [48–50]. Several studies have reported an increase in network unbalance with the inclusion of single-phase EVs [50–52]. Moreover, the study in [53] found that

the thermal degradation of cables would be accelerated by an increased number of EVs charging during the evening peak.

Overall, there is compelling evidence for the need for increased flexibility in distribution networks, both in terms of the management of customer-connected assets, as well as within the network components. A DNO in the UK recently published [54] their strategy to actively manage the output of generation assets connected by customers in order to avoid breaching network constraints, a clear step in the evolution towards a Distribution System Operator (DSO) role. Instead of offering a traditional firm connection, the customer agrees to have a flexible connection, which allows the operator to regulate the output of its generation assets, as a form of Active Network Management (ANM). This reduces both the decision time and the cost of these connections. Another recent project investigated the benefits of using automatic switches to mesh LV substations in order to balance the loading between the transformers and release latent capacity [55]. Similarly, a DNO in the UK is set to trial the benefits of replacing conventional transformers with an alternative based on power electronics to provide real-time voltage regulation and power flow management, further increasing the flexibility [56].

## 1.5 Power Electronics Devices as Enablers of Flexibility

Distribution networks are normally operated radially, thus a straightforward strategy to limit the impact of LCTs on network assets and to provide additional flexibility is to mesh the feeders and lines through switching. The term *flexibility* is used here to refer to the DNO's ability to control the network assets in order to create additional paths for the power flows. This can help distribute the power from the loads and generators more uniformly across the transformers and feeders, allowing the DNO to fully utilise the network capacity.

The authors of [57] analysed how a simple meshed structure could affect the parameters of a distribution network. In particular, the authors analysed how it influenced the amount of allowable DG in the network. One important point highlighted in the study was that optimising both the number and location of the ties, according to the loading and generation conditions in the network, is essential to reap the benefits of the meshed structure, to avoid worsening the loading conditions of weak feeders, and to minimise costs. A more realistic examination of the benefits of meshing LV feeders to increase the amount of low-carbon technologies that can be installed in distribution networks was carried out in [58]. Three factors were considered and compared to assess the performance of the meshed operation versus the standard radial configuration: voltage problems, thermal issues, and energy losses. The results showed that the number of feeders presenting problems was greatly reduced when the network was meshed and, hence, the allowable penetration of low-carbon technologies was increased. Nonetheless, it is worth noting that in the previous two studies the

feeders were simply meshed at their end points by closing a switch. Consequently, the voltage and network impedance governed the power flow, and the increase in the hosting capacity of the most problematic feeders came at the expense of the neighbouring feeders, since they shared the load/generation. It is also worth highlighting that one of the feeder's parameters can act as a bottleneck of its hosting capacity. Moreover, a direct connection between feeders and substations is not always viable due to the existing network boundaries, which present differences in the voltage and/or phase. A simple tie can lead to circulating currents and protection issues, as well as excessive fault levels [59, 60].

To avoid a redesign of the protection system, derived from a solid connection between feeders, and to provide extra controllability, power electronic converters could be used as the main interface connecting feeders and substations together. Unlike traditional meshing, the current injections by power converters can be adjusted to a desired value almost instantaneously, hence they can provide an additional level of network flexibility. Also, because of their fast operation, they do not contribute to the fault current in the network [61]. Moreover, they have the additional advantage of providing a degree of electrical isolation when connected back-to-back (i.e. AC-DC-AC). Thus a fault on one of the AC sides would not pass through the DC bus to the other AC side, which can potentially allow connections that cross network boundaries [62]. For instance, they are able to mesh substations at different voltage levels, such as 11 kV and 6 kV, as well as connecting feeders supplied from different transformers [62]. In addition, power converters are able to control the power flow irrespective of the topology, whereas in the case of a solid connection the impedance of the network governs the power transfers. Therefore, power converters offer the possibility of regulating the power flow between the feeders and between substations, effectively pooling the demand/generation coming from LCTs and optimising the hosting capacity of the whole network. However, to provide such flexibility power electronic converters must be accompanied by a suitable control scheme, including both low-level and high-level controllers. For the case of an ANM scheme, a high-level control is needed to calculate the set points for each power converter, whilst taking into account the overall network characteristics. Designing a control suitable for real-time operation that acts optimally whilst accounting for the network constraints and the constant changes in the load/generation can be complex. This problem will be investigated in more detail in the following chapters.

## 1.6 Research Aims

This chapter has reviewed some of the challenges faced by DNOs as a result of the increased volatility in the power flows brought about by the integration of large volumes of LCTs into distribution networks. Additionally, it has highlighted the need for flexibility, both from the customer side and within the physical components of the network (e.g., enabled by switches or power electronic converters). From the literature

reviewed, it is clear that as the adoption of LCTs accelerates, a more comprehensive form of ANM including active management of the power flows will become necessary for DNOs, as they transition towards a DSO role. LV networks should no longer be considered passive, where components are installed on a fit-and-forget basis. Instead, their operation and control should be dynamic and responsive to the constant changes in load and generation, to achieve maximum operational efficiency. However, designing control schemes that incorporate realistic models of LV networks can be challenging, due to their intrinsic characteristics and the constant power fluctuations. In this context, this thesis investigates how the operation of LV networks can be enhanced by implementing an ANM scheme using power converters, which allow complete control of the power flows. The main focus of this thesis is on the development and evaluation of a high-level control scheme that would produce optimal power set points for embedded power electronic converters for a given time period. Several options are investigated, ranging from traditional nonlinear optimisation to supervised learning. Reinforcement learning is also analysed as an alternative optimisation framework. These options are compared in terms of computational efficiency, visibility requirements and generalisation ability. Emphasis is given to the incorporation of realistic models of distribution networks into the control design, to accurately reflect their behaviour and to ensure their operation is improved. In addition, practical constraints, such as the limited availability of monitoring devices, are taken into account in the control design to ensure an optimal operation. Ultimately, the proposed controller could enable an ANM scheme based on power electronic converters to be implemented in real-time operation.

## 1.7 Thesis Structure

The remainder of this thesis is structured as follows:

- **Chapter 2** introduces the framework used to model realistic distribution networks – in particular, LV networks – for power flow and optimal power flow (OPF) studies. It draws attention to the differences between transmission and distribution networks and emphasises the accurate representation of the phase imbalance and the lack of load diversity present in LV networks. This modelling framework forms the basis for the subsequent chapters.

- **Chapter 3** describes how power converters can be integrated into the optimal power flow formulation to model the active management of power flows in distribution networks. The proposed formulation considers the connection of pairs of converters back-to-back and also the connection of several converters in a multi-terminal fashion via DC connections. Case studies are presented to demonstrate the benefits of the proposed active network management scheme in comparison to the more traditional network meshing through switching. Finally, a discussion is presented on the suitability of the

proposed OPF as a high-level control for real-time operation.

- **Chapter 4** proposes an alternative control method based on supervised learning to improve the computational efficiency of the OPF for real-time operation control. A neural network is trained to map optimal set points, obtained through an OPF, to available network measurements using regression.

- **Chapter 5** presents an alternative optimisation framework based on reinforcement learning and neural networks and compares two different approaches, model-free and model-based. The model-free algorithm makes use of a *black-box* training environment, representing a model of the LV network, to train a neural network to choose optimal actions. The model-based alternative incorporates the power flow equations as a single computational graph, which facilitates the calculation of the gradients via automatic differentiation and enables the use of gradient-based optimisation algorithms. The advantages of both these approaches over traditional optimisation and over supervised learning are evaluated through case studies.

- **Chapter 6** summarises the research findings and provides conclusions and recommendations. It also discusses the assumptions made in the study, its limitations, and the areas that fall outside of its scope. It then suggests different avenues that can be considered for future research to address those areas.

Note that detailed review of literature regarding specific optimisation algorithms and machine learning techniques will be conducted in the respective chapters where those techniques are used.

## 1.8 Contribution

To improve the flexibility of distribution networks, this thesis proposes using ANM schemes to optimally distribute the power flows in real-time using power electronic converters. Because of the challenges found during real-time operation, such as the unavailability of network monitoring in distribution networks and the computational efficiency required, algorithms based on supervised learning and reinforcement learning are proposed as high-level controllers for these ANM schemes.

Although supervised learning can be straightforward to implement, the results show that it has low accuracy and most importantly it cannot generalise to out-of-sample data. Reinforcement learning is proposed as an alternative optimisation framework for the ANM problem to overcome those challenges. Two main algorithms based on reinforcement learning are tested: 1) a model-free algorithm, which uses a *black-box* model of the distribution network to train a neural network based on continuous interactions, and 2) a model-based algorithm that uses the power flow equations explicitly, along with gradient information, to train a neural

network to take optimal decisions. The advantages of both algorithms, compared to traditional optimisation, are demonstrated through case studies, using realistic models of distribution networks.

The results show that the proposed algorithms based on reinforcement learning can take advantage of automatic differentiation software and gradient-based optimisation techniques to keep a high computational efficiency whilst achieving accurate minimisation of the chosen objective, with a difference of less than 10% compared to the OPF. Moreover, the trained neural networks can generalise to out-of-sample loading conditions, whilst its recall stage is very fast, in the range of milliseconds for each time step. Also, unlike traditional optimisation, once trained the proposed controller can operate under partial observability, requiring only a subset of network measurements as input.

The end goal of this thesis is to provide a comprehensive assessment of these alternative control methods, in order to inform the deployment of ANM schemes based on power electronic converters in the near future.



# Chapter 2

# Distribution Network Modelling

## 2.1 Distribution Networks

Before the major improvements in energy storage devices of recent years [3], it was extremely expensive to store large volumes of electrical energy. Without storage, the balance between supply and demand in the power system had to be maintained at all times. Such time-critical balancing was made possible by the inter-linked structure of the system and the use of sophisticated control systems, such as the automatic generation control (AGC) found in large generators [63]. A power system is usually structured vertically. Power is generated by large power stations, transported at extra high voltage (EHV) via long transmission lines and delivered at high voltage (HV) to load centres via multiple substations and feeders as part of a distribution network [64]. Fig. 2.1 provides an illustration of the power system structure, showing voltage levels that are typical for the UK.

Distribution networks in the UK are usually fed from the transmission network via Grid Supply Points (GSPs), operating at 132 kV. Each GSP supplies a number of 33 kV Bulk Supply Points (BSPs) with several primary substations downstream. Substations can be seen as the point of interconnection between lines at different voltage levels through transformers. They also house the switchgear and protection equipment, along with measurement instrumentation such as Supervisory Control and Data Acquisition (SCADA) systems [66]. Primary substations are normally operated at 33/11 kV and have several outgoing feeders that supply a given number of secondary substations. These low-voltage substations are operated at 11/0.4 kV, with domestic and commercial customers being supplied at the lowest voltage level of 230 V [30, 67, 68]. This is illustrated more clearly in Fig. 2.2, where the connection to the transmission network is illustrated as an equivalent voltage source. The "last mile" of the distribution system is known as the low-voltage (LV) network and is the most extensive and complex part of the grid, given its diverse topology, proximity to the customers and large number of individual nodes [25].

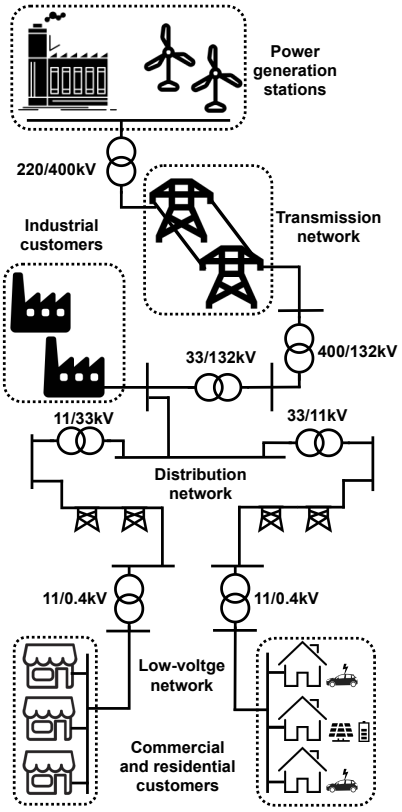



Figure 2.1: Typical power system structure with standard voltage levels for the UK. Large generation stations usually include several voltage transformation stages before connecting to the transmission network, for instance from 33 kV or 66 kV to 220 kV in the case of offshore windfarms [65].

According to Ofgem, the total length of underground LV cables is approximately 328,038 km, compared to 3,190 km for cables at 132 kV [37]. Ground-mounted LV substations supply on average 3 to 4 feeders, with a higher number for urban areas. Although several configurations can be adopted for LV networks, such as ring or mesh, they are normally operated radially [66, 67]. Some of the feeders might be interconnected so that, in the case of faults, the supply can be restored from a different point. However, the switches that achieve this are considered to be Normally Open Points (NOPs), as shown in Fig. 2.2 [66].

To model and analyse a given distribution network, equivalent circuits are used for the different components, using mainly underground cables (for urban and suburban settings), overhead lines (for rural areas) and transformers [66, 67]. Lower voltage cables and overhead lines have their conductors spaced closer together than higher voltage cables and lines. Closer spacing gives rise to less flux-linkage and self-inductance which in turn means low-voltage lines have lower $X : R$ ratio [67]. Generally the conductor cross-sections are broadly similar and so resistance changes less with voltage level. As a result of the lower $X : R$ ratio, the effect of the reactive power on the voltage magnitude is small but, on the other hand, the impact of active power injections on the voltage profile is greater and large voltage fluctuations can occur [30, 67]. This is of particular importance given the large and intermittent power injections from LCTs. During peak load conditions, large voltage drops might often be observed in the middle or the end of the feeder.

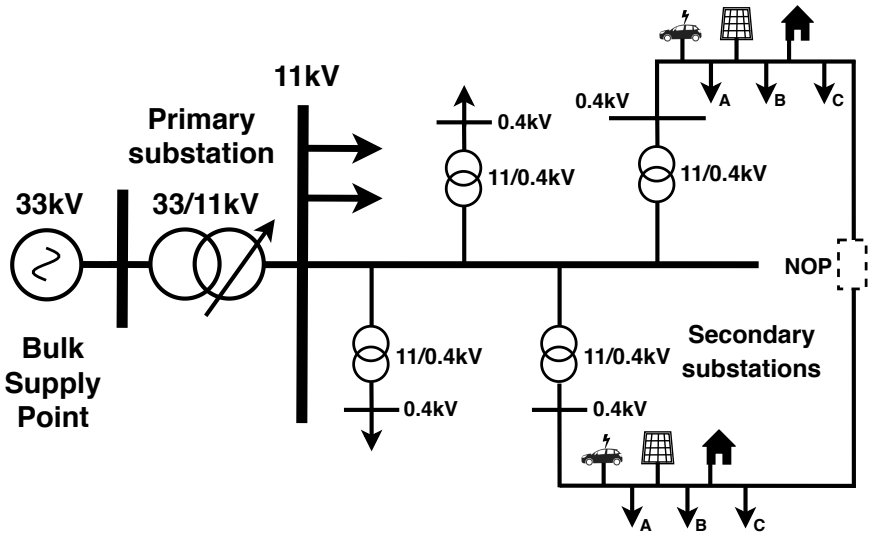



Figure 2.2: Distribution network structure with typical voltage levels for the UK [69]. The NOP can be used to connect two LV substations, via one of their feeders, to restore the supply in case of a fault.

At times of peak generation from solar PV installations and low loading conditions, reverse power flow and increased voltage magnitude above the nominal may occur. The shunt admittance caused by the capacitance of a cable or line is normally ignored, since its effect is minimal in LV networks, due to the short length of the lines compared to transmission lines [39,70,71]. Typical positive-sequence ($R_1$, $X_1$) and zero-sequence ($R_0$, $X_0$) impedances per unit length ($\Omega$/km), and ampacity for cables used in LV networks in the UK are shown in Table 2.1 [67,69,72,73].

Table 2.1: Typical LV cables data.

| Code | $R_1$ | $X_1$ | $R_0$ | $X_0$ | Ampacity [A] | Description |
|---|---|---|---|---|---|---|
| 4c 300 | 0.102 | 0.072 | 0.291 | 0.087 | 465 | 4 cores 300mm$^2$ |
| 4c 185 | 0.166 | 0.073 | 0.412 | 0.094 | 355 | 4 cores 185mm$^2$ |
| 4c 120 | 0.253 | 0.068 | 0.885 | 0.080 | 280 | 4 cores 120mm$^2$ |
| 4c 95 | 0.322 | 0.074 | 0.805 | 0.092 | 245 | 4 cores 95mm$^2$ |
| 4c 70 | 0.446 | 0.071 | 2.540 | 0.082 | 205 | 4 cores 70mm$^2$ |

Transposing the conductors by swapping their positions at regular intervals is normally done in overhead transmission lines to keep the impedances of the phase conductors symmetric, even when they are not spaced equilaterally [64]. However, this is not the case in distribution networks where a major part is composed of underground cables. Thus distribution lines are untransposed. This results in different self and mutual impedances between the phases, and thus different voltage drops in each phase, even if the supply and the loads are balanced [39]. On top of this, the majority of the customers in European networks are normally connected to a single phase and distributed unevenly across phases [74–76]. In combination with the different power demands of customers through the day, a large difference in the line currents between

phases is often observed, with one of the phases highly loaded compared to the others [67].

It is normal practice for the cores of the three phase cores to have a larger cross-sectional area than the neutral [77]. Given the inverse relationship between cross-sectional area and power losses, this means that having current flowing in the neutral results in significantly more losses than when the currents are balanced. Because of the various features of networks and customer connections just discussed, it is unreasonable to assume a balanced three-phase operation for distribution networks and so reduction to a single-phase equivalent model is unjustified [39]. It is thus imperative to model each of the phases of the network components [39, 78].

In a network model, transformer connections are modelled using a nodal admittance matrix similar to the one calculated for the lines, constructed from the shunt and series impedances of the transformer. These quantities depend on parameters, such as internal power losses, apparent power rating and voltage level. The losses are usually divided into no-load losses (i.e., those due to hysteresis and eddy-current in the iron core which arise even without power flowing) and load losses (i.e., those which depend on the current passing through the windings and also known as copper losses) [67, 75]. Parameters of typical transformers used in distribution networks in the UK are included in Table 2.2 [69, 79, 80]. The impedance components are provided in per-unit, that is as a fraction of the impedance base, calculated using the transformer power rating in MVA ($S_{base}$) and its voltage base ($V_{base}$) in kV [64], as in (2.1).

$$Z_{base} = \frac{(V_{base})^2}{S_{base}} \tag{2.1}$$

Usually, a transformer can have a loading above its nameplate rating for a short period of time without affecting its lifespan. This is due to the cyclic nature of the loads and the fact that transformers rarely experience continuous loading over a 24-hour period [81]. The percentage of overload that is acceptable for a given transformer depends on the initial demand, the ambient temperature and the duration of the overload [81].

Table 2.2: Typical parameters for transformers used in LV networks [69, 79].

| **Sector** | **Rating [kVA]** | **Load losses at 75°C [W]** | **No-load losses [W]** | **Impedance Z [%]** | **$R_{LV}$ [pu]** | **$X_{LV}$ [pu]** |
|---|---|---|---|---|---|---|
| Highly urban | 1500 | 15810 | 1400 | 5.0 | 0.01054 | 0.0489 |
| Highly urban | 1000 | 11000 | 1350 | 4.75 | 0.0110 | 0.0462 |
| Highly urban | 800 | 7410 | 1000 | 4.75 | 0.0093 | 0.0466 |
| Urban | 500 | 5100 | 680 | 4.75 | 0.0102 | 0.0464 |
| Urban | 315 | 3420 | 580 | 4.75 | 0.0109 | 0.04624 |
| Sub-urban | 200 | 2900 | 540 | 4.75 | 0.0150 | 0.0450 |
| Rural | 100 | 1750 | 320 | 4.5 | 0.0175 | 0.0415 |
| Rural | 50 | 1100 | 190 | 4.5 | 0.0219 | 0.0393 |

## 2.2 Power Flow

The power flow algorithm was initially developed to analyse power systems from the point of view of the upstream generation and transmission network, to support the growing demand for electricity in the preceding decades [82]. A power flow calculation is used to provide a snapshot of the steady-state operation of a network, or part of it, which includes generators, transmission or distribution lines and loads [64, 83, 84]. The output of the calculation provides the voltage profile along the buses and feeders, and the currents flowing through the lines. Such analysis would be part of the planning process and also used during operation [63, 64]. Power flow studies are usually carried out when considering network expansions and reconfigurations, new loading conditions or when incorporating new equipment [64]. In distribution networks, power flow studies are used in the operation and planning stage, when considering future network requirements or to verify network reconfiguration arrangements [67].

In its AC form, the power flow is a non-linear problem and thus it requires iterative methods, such as the Gauss-Siedel or the Newton-Raphson methods [63, 83]. The standard Newton-Raphson method, and its decoupled variations, were conceived with the characteristics of transmission networks in mind. Thus, they tend to perform poorly for radial distribution networks or they ignore important characteristics from distribution networks [66]. For this reason, neither the Gauss-Siedel nor the Newton-Raphson method will be discussed in this thesis. The interested reader is directed to the abundant literature on more traditional methods for transmission networks [64, 66].

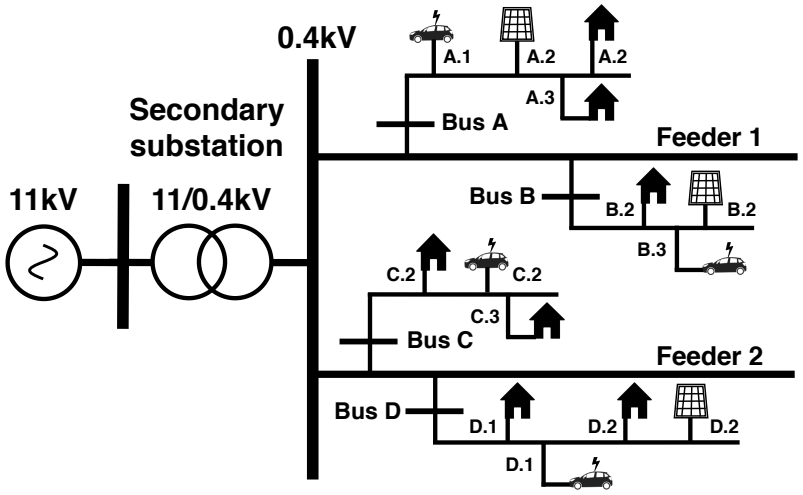



Figure 2.3: Diagram of two LV feeders supplied from the same primary substation, with typical voltage levels for the UK. The usual single-phase connection of LCTs is illustrated.

Power flow calculations in transmission networks normally make the assumption of three-phase balance, which for the reasons discussed in section 2.1, and because of the large load diversity at higher voltage levels, is reasonable. This makes it possible to analyse only one phase and extend the results to the other two

phases. However, due to the physical differences in distribution networks, and because of the large number of single-phase loads at each bus, as described in section 2.1, the traditional power flow formulation is not suitable to model these networks accurately. The diversity of the loads decreases in the lower voltage levels of distribution networks, where there are considerable differences in demand between consumers in each of the phases and at different times [85–87]. The same is true for LCTs connected in LV networks, as was shown in section 1.3. This is illustrated in Fig. 2.3, which depicts a primary substation supplying two LV feeders, each of them with its respective loads and LCTs connected to a bus and a specific node (e.g., A.1, B.2, etc.). Detailed models of the transformers, feeders and loads are thus required to take into account the high-phase imbalance and low $X : R$ ratio of the cables [39].

For a given three-phase line, the first step is to calculate its primitive impedance matrix $Z_{prim}$ using its conductor data and the modified Carson equations. Although the primitive matrices do not provide any information about how the different components are connected to each other, they model each individual component and, when combined together, they can model a given multi-phase network. For overhead and underground lines with three conductors (one conductor per phase) and a neutral, the primitive self-impedance $Z_{ii}$ and mutual-impedance $Z_{ij}$ of the conductors $i$ and $j$ can be obtained via the modified Carson's equations (2.2) and (2.3), respectively [39]:

$$Z_{ii} = r_i + 0.00158836 \cdot f + j0.00202237 \cdot f(\ln \frac{1}{\text{GMR}_i} + 7.6786 + \frac{1}{2} \ln \frac{\rho}{f}) \tag{2.2}$$

$$Z_{ij} = 0.00158836 \cdot f + j0.00202237 \cdot f(\ln \frac{1}{D_{ij}} + 7.6786 + \frac{1}{2} \ln \frac{\rho}{f}) \tag{2.3}$$

where:

$f$ = frequency in Hz
$\rho$ = earth resistivity in Ohm-m
$r_i$ = resistance of conductor $i$
$\text{GMR}_i$ = geometric mean radius of conductor $i$
$D_{ij}$ = spacing between conductors $i$ and $j$

Assuming a 60Hz frequency and an earth resistivity of 100 Ω-m, the primitive self-impedance $Z_{ii}$ and mutual-impedance $Z_{ij}$ in Ω/mile can be obtained as in (2.4) and (2.5), respectively [39]:

$$Z_{ii} = r_i + 0.09530 + j0.12134(\ln \frac{1}{\text{GMR}_i} + 7.93402)\ \Omega/\text{mile} \tag{2.4}$$

$$Z_{ij} = 0.09530 + j0.12134(\ln \frac{1}{D_{ij}} + 7.93402)\ \Omega/\text{mile} \tag{2.5}$$

The primitive impedance $Z_{prim}$ is then obtained as in (2.6) [39]:

$$\left[Z_{prim}\right] = \begin{bmatrix} Z_{ii} & Z_{ij} & Z_{in} \\ Z_{ji} & Z_{jj} & Z_{jn} \\ Z_{ni} & Z_{nj} & Z_{nn} \end{bmatrix} \tag{2.6}$$

where:

$Z_{in}$ = mutual-impedance between conductor $i$ and the neutral $n$
$Z_{jn}$ = mutual-impedance between conductor $j$ and the neutral $n$
$Z_{nn}$ = self-impedance of the neutral $n$

Then, the phase impedance matrix $Z_{abc}$ is obtained by applying the Kron reduction to the primitive impedance matrix, as in (2.7) and (2.8) in matrix format [39]. For the subsequent analysis, all four-wire lines will be reduced to their three-wire representation using the Kron reduction to calculate the impedance matrix [39].This reduction assumes that the neutral of the line is multi-grounded [39].

$$Z_{abc} = Z_{ij} - \frac{Z_{in}}{Z_{nn}} Z_{nj} \tag{2.7}$$

$$\left[Z_{abc}\right] = \begin{bmatrix} Z_{aa} & Z_{ab} & Z_{ac} \\ Z_{ba} & Z_{bb} & Z_{bc} \\ Z_{ca} & Z_{cb} & Z_{cc} \end{bmatrix} \tag{2.8}$$

The symmetrical components framework proposed by Fortescue can be used to represent an unbalanced multi-phase system of phasors as three sets of components: positive-, negative- and zero-sequence [64,88,89]. The phase impedance can be expressed as its equivalent sequence impedance $Z_{012}$ by using the transformation in (2.9) [39]:

$$Z_{012} = [A_s]^{-1} Z_{abc} [A_s] \tag{2.9}$$

where:

$$A_s = \begin{bmatrix} 1 & 1 & 1 \\ 1 & a^2 & a \\ 1 & a & a^2 \end{bmatrix} \tag{2.10}$$

and $a_s$ is the phasor rotation operator:

$$a_s = 1\angle 120^\circ \tag{2.11}$$

The sequence impedance can be expanded as in (2.12):

$$\left[Z_{012}\right] = \begin{bmatrix} Z_0 & 0 & 0 \\ 0 & Z_1 & 0 \\ 0 & 0 & Z_2 \end{bmatrix} \tag{2.12}$$

where:

$Z_0 =$ is the zero sequence impedance

$Z_1 =$ is the positive sequence impedance

$Z_2 =$ is the negative sequence impedance

Given this transformation, and its inverse in (2.14), it is possible to obtain the approximate phase impedance of a line model only from its zero and positive sequence impedance, as those in Table 2.1, if the complete phase impedance is not available. This is typically the case in some distribution networks [90].

$$Z_{abc} = [A_s] Z_{012} [A_s]^{-1} \tag{2.13}$$

$$\left[Z_{abc}\right] = \frac{1}{3} \begin{bmatrix} 2Z_1 + Z_0 & Z_0 - Z_1 & Z_0 - Z_1 \\ Z_0 - Z_1 & 2Z_1 + Z_0 & Z_0 - Z_1 \\ Z_0 - Z_1 & Z_0 - Z_1 & 2Z_1 + Z_0 \end{bmatrix} \tag{2.14}$$

Although this inverse transformation assumes transposition [39], its use is subject to the cable data being available to build the network model. In the case of transformers, their impedance values are usually obtained from short-circuit and open-circuit tests [63]. A numerical example of the process followed to construct the phase impedance for a line can be found in **Appendix A**.

Normally, the known quantities in distribution networks are the voltage at the substation (kept close to 1.0 pu), the active and reactive power requested by the loads, and an estimate of the load type (e.g. constant power, constant current, etc.) [39]. By using the parameters of the network components such as lines and transformers, the power flow can determine: (1) the voltage magnitude and angle at all the nodes and (2) the current magnitude and angle flowing in each line. These two results can then be used to calculate the total power fed from the substation, the power losses in the transformers and lines and the power absorbed by the loads, depending on their type. Also, the voltage drop along the feeder can be calculated to see if compensation is required. According to the British Standard EN50160:2010+A3:2019 [91], under normal operating conditions in the UK, the limits for voltage variations at LV networks are set at $\pm$10% of the nominal voltage (230 V for single-phase connected customers). This standard also states that the voltage unbalance should not exceed more than 2% of the positive sequence component. The voltage unbalance factor (VUF), or its percentage can be calculated using the ratio between the magnitudes of the negative ($V_2$) and the positive

sequence ($V_1$) components of the voltage of each three-phase node, as in (2.15) [92, 93].

$$\%\text{VUF} = \frac{|V_2|}{|V_1|} \cdot 100 \tag{2.15}$$

The negative ($V_2$) and positive ($V_1$) sequence components are defined as in (2.16) and (2.17), respectively:

$$V_2 = \frac{1}{3}(V_a + a_s^2 V_b + a_s V_c) \tag{2.16}$$

$$V_1 = \frac{1}{3}(V_a + a_s V_b + a_s^2 V_c) \tag{2.17}$$

where $V_a$, $V_b$ and $V_c$ are the voltages at each of the phases.

Depending on the behaviour of the loads with respect to the voltage at their terminals, they can be modelled as: (1) constant power, (2) constant current, (3) constant impedance or (4) a combination of the above. The ZIP model [94] can be used to model loads that show a mixed behaviour, comprising the first three types above: constant impedance (Z), constant current (I) and constant power (P).

## 2.3 Fixed-point Algorithm for Power Flow

The power flow problem for distribution networks can be formulated by constructing the nodal admittance matrix ($[Y]$) of a network. This nodal admittance matrix ($[Y]$), also known as the system admittance matrix ($[Y_{system}]$), is built by combining the primitive admittance matrices ($[Y_{prim}]$) of each element in the network, such as transformers, lines and loads. This formulation is used in OpenDSS, an open-source simulator for distribution networks [95], to perform power flow calculations. These $[Y_{prim}]$ matrices depend on the nature of the element modelled, either a voltage source, line, load or generator. Each of these elements is associated to a node or a pair of nodes within a bus, for instance a line connected between node A and B or a load connected to node A. Fig. 2.4 illustrates the concept of bus and node more clearly: where a line is connecting bus A with bus B and internally, it is connected to all phases or nodes (A.1, A.2 and A.3), including the neutral (A.n). However, another line makes a single phase connection between bus B and bus C. This naming convention will be used for the rest of the thesis.

Once the admittance matrix of the system ($[Y_{system}]$) is constructed, it can then be used as part of a system of equations that represent the nodes voltages ($[V_{nodes}]$), or simply ($[V]$), and all the injected currents ($[I_{nodes}]$), or simply ($[I]$) [95]. The nodal admittance matrix relates injected currents to nodes voltages as in (2.18) and in short (2.19).

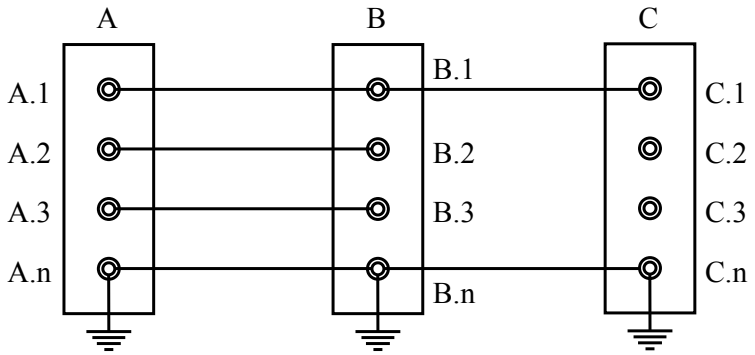



Figure 2.4: Line connections between bus A and bus B and between bus B and bus C, depicting connections between internal nodes [69].

$$\begin{bmatrix} I_1 \\ I_2 \\ \vdots \\ I_n \end{bmatrix} = \begin{bmatrix} Y_{1,1} & Y_{1,2} & \dots Y_{1,n} \\ Y_{2,1} & Y_{2,2} & \dots Y_{2,n} \\ \vdots & \vdots & \vdots \\ Y_{n,1} & Y_{n,2} & \dots Y_{n,n} \end{bmatrix} \begin{bmatrix} V_1 \\ V_2 \\ \vdots \\ V_n \end{bmatrix} \tag{2.18}$$

$$[I] = [Y][V] \tag{2.19}$$

$$[S] = [V][I]^* \tag{2.20}$$

Lines and transformers can be fully represented using their primitive admittance matrix $[Y_{prim}]$, thus their current can be calculated as in (2.19) and the complex power as in (2.20). Loads and generators, on the other hand, can have non-linear behaviour depending on their specific type. For instance, loads that absorb a constant amount of power will have their current vary inversely, in proportion to the voltage at its terminals. The loads and generators which exhibit a nonlinear behaviour are represented using a Norton equivalent, where the parallel admittance matrix is obtained using the nominal voltage and incorporated into the admittance matrix of the system $[Y_{system}]$. The non-linear behaviour between the current and the voltage is represented by a compensation current and requires several iterations of $k$ on the voltage. The main current vector $[I]$ in (2.18) is populated by the compensation currents of the loads and generators, as these are the only elements that inject current into the network. For a constant power load connected to node $i$, its compensation current $I_{comp,load}$ can be calculated as in (2.21):

$$I_{comp,load} = I_{linear,load} - I_{inj,load} \tag{2.21}$$

where the linear part $I_{linear,load}$ depends only on its primitive admittance matrix and the voltage at the node

$i$ to which it is connected, as in (2.22):

$$[I_{linear,load}] = [Y_{prim}] \times [V_i^k] \tag{2.22}$$

For daily power flow calculations the $Y_{prim}$ of the loads and generators does not change as it only depends on the peak active and reactive power for the load/generator and their respective nominal voltage. The current absorbed by the load, $I_{inj,load}$, depends on the power requested and the voltage at the load terminals. Thus it can be calculate as the conjugate of the complex power ($S_i$) over the voltage at node $i$, as in (2.23)

$$I_{inj,load} = \left(\frac{S_i}{V_i^k}\right)^* \tag{2.23}$$

where the complex power ($S_i$) can be decomposed into active ($P_i$) and reactive power ($Q_i$), as in (2.24):

$$S_i = P_i + jQ_i \tag{2.24}$$

A numerical example of the calculation of $Y_{prim}$ and the compensation current for loads and generators is given in **Appendix A**.

Given this non-linear relationship, an iterative method is needed to update the compensation currents at each iteration with the new voltage estimate. A method that has been proven to be efficient for this purpose is the use of fixed-point iterations [71, 95]. This is one of the two algorithms used by OpenDSS to calculate the power flows for large multi-phase networks [95]. By making an initial guess $k$ at the voltage ($V_i^k$), the compensation current of a constant power load can be calculated as in (2.25):

$$I_{comp,load} = [Y_{prim}] \times [V_i^k] - \left(\frac{P_i + jQ_i}{V_i^k}\right)^* \tag{2.25}$$

Once the $[Y_{system}]$ is built and all the compensation currents are calculated and incorporated into $[I_{nodes}]$, (2.19) can be solved to find the voltages at all the nodes ($[V]$) using fixed-point iterations [95]. One of the advantages of this is that the $[Y_{system}]$ remains constant during the fixed-point iterations, unless the physical structure of the network changes due to switching. Using the $[Y_{system}]$, the voltages at all nodes in the network can be calculated recursively as in (2.26)

$$[V_{nodes}^k] \leftarrow [Y_{system}]^{-1} \times [I_{nodes}(V_{nodes}^{k-1})] \tag{2.26}$$

where $[I_{nodes}]$ is a vector containing all the currents injected by the loads and any generators connected to the network and it depends on the nonlinear behaviour of the respective element. If this converges as $k \to \infty$

then it achieves $[V_{nodes}^{k+1}] = [V_{nodes}^{k}]$, or simply $[V_{nodes}^{k}] = [V_{nodes}]$, hence it can be written as in (2.27):

$$[V_{nodes}] = [Y_{system}]^{-1} \times [I_{nodes}(V_{nodes})] \tag{2.27}$$

which is equivalent to equation (2.19), as required.

The voltages at all the nodes are calculated at every iteration and are used to update the compensation current of non-linear elements (e.g. constant power loads). The algorithm stops until the error between the previous and the present iteration is below a certain tolerance ($10^{-5}$ is the default used in OpenDSS [95]). The voltage error for node $i$ can be calculated as in (2.28):

$$V_{error,i} = \frac{\left|\left|V_i^{(k)}\right| - \left|V_i^{(k-1)}\right|\right|}{V_{base}} \tag{2.28}$$

where the $V_{base}$ is the base voltage. The full fixed-point iteration algorithm for the calculation of the power flow in a multi-phase network is described in Algorithm 1 [95].

**Algorithm 1** Fixed-point algorithm for power flow

Get primitive admittance matrices ($Y_{prim}$) of each element in the network
Get nodal admittance matrix of the system ($Y_{system}$)
Get voltage of the equivalent source at the substation assuming nominal voltage $V_{base}$ and balance
Get injection current from the substation $I_{source}$
1: Start iterations $k = 0$
2: Initial guess at the voltage vector $V_{nodes}^{(k-1)}$ without the loads compensation currents as in (2.26)
3: **while** $V_{error} \geq 10^{-5}$ **do**
4: Update the iterator counter $k = k + 1$
5: Calculate linear part of current for the loads as in (2.22)
6: Calculate injection current for the loads as in (2.23)
7: Calculate compensation current for the loads as in (2.25)
8: Update the voltage vector $V_{nodes}^{(k)}$ including the compensation current of the loads as in (2.26)
9: Calculate the $V_{error}$ as in (2.28)
10: **end while**
11: Calculate the current in all lines with the final $V_{nodes}$ and the respective $Y_{prim}$ as in (2.19)
12: Calculate the power flow in all lines with the final $V_{nodes}$ and the respective currents as in (2.20)

This fixed-point iteration algorithm is used in OpenDSS [95] for power flow and has been shown to perform well in realistic distribution networks [58, 96].OpenDSS has several modes to support the solution of power flows when using load profiles, however the most common ones are snapshot, daily and yearly. A more detailed description of the power flow formulation and the fixed point algorithm is given in **Appendix A**.

## 2.4 Optimal Power Flow

Unlike the power flow algorithm, which calculates the voltages at all the nodes given a set of network conditions, the goal of the optimal power flow algorithm (OPF) is to calculate the *optimal* actions by control devices that result in, typically, a minimisation of a cost-function to a level that is feasible given the network physical constraints [97]. The cost-function could be the cost of running the generators or could reflect some desirable network operating feature [98]. For instance, the OPF has traditionally been used to calculate optimal active and reactive power injections from generators, subject to the transmission network constraints. The most common objective function is to minimise the production cost of the generators. The network constraints are represented by equalities and inequalities. The active/reactive power flow equations of the network are normally represented by equality constraints, whereas the inequality constraints incorporate any limits on the generators, the buses, or the lines. The traditional OPF can be formulated as in (2.29) [99]:

$$\underset{x}{\text{minimize}} \quad f(x) = \sum_{i=1}^{N} f_i(P_{g,i}) \tag{2.29a}$$

$$\text{subject to} \quad P_{g,i} - P_{d,i} - P_i(V,\theta) = 0, \tag{2.29b}$$

$$Q_{g,i} - Q_{d,i} - Q_i(V,\theta) = 0, \tag{2.29c}$$

$$P_{g,i}^{min} \leq P_{g,i} \leq P_{g,i}^{max}, \tag{2.29d}$$

$$Q_{g,i}^{min} \leq Q_{g,i} \leq Q_{g,i}^{max}, \tag{2.29e}$$

$$V_i^{min} \leq V_i \leq V_i^{max}, \tag{2.29f}$$

$$P_{l,m} \leq P_{l,m}^{max} \tag{2.29g}$$

where:

$P_{g,i}, Q_{g,i}$ = are the active and reactive power injections of the generator at node $i$

$P_{d,i}, Q_{d,i}$ = are the active and reactive power injections of the load at node $i$

$P_{l,m}$ = is the power limit at line $m$

$V_i$ = is the voltage at node $i$

The cost function usually is represented by a quadratic function that depends on the power generated as in (2.30):

$$f_i(P_{g,i}) = a + bP_{g,i} + cP_{g,i}^2 \tag{2.30}$$

where here $a$, $b$ and $c$ are constants associated with the fuel or other generator parameters.

The power flow equations can be represented in their polar form as in (2.31) and (2.32). Other forms such as the rectangular coordinates for the voltage can also be used [100, 101].

$$P_i(V,\theta) = V_i \sum_{j=1}^{N} V_j[G_{ij}\cos(\theta_i - \theta_j) + B_{ij}\sin(\theta_i - \theta_j)] \tag{2.31}$$

$$Q_i(V,\theta) = V_i \sum_{j=1}^{N} V_j[G_{ij}\sin(\theta_i - \theta_j) - B_{ij}\cos(\theta_i - \theta_j)] \tag{2.32}$$

where:

$V_i, V_j$ = are the voltages at node $i$ and $j$

$G_{i,j}, B_{i,j}$ = are the conductance and the susceptance of a line between nodes $i$ and $j$

$\theta_i, \theta_j$ = are the voltage angles at node $i$ and $j$

In its typical form, the OPF is a non-linear, non-convex optimisation problem that can be solved through several methods, such as the gradient method and sequential quadratic programming. [102]. The intrinsic non-convexity of the OPF problem makes it computationally challenging for large networks. The power injections in (2.31) and (2.32), as well as the voltage limit constraint in (2.29) are the main sources of non-convexity [101, 103] and they render the OPF NP-hard [104]. However, one method that has gained traction in recent years, due to the approach it uses to manage the inequality constraints and its rapid convergence, is the primal-dual Interior Point Method (IPM) [99, 100, 105].

In the study of distribution network control described in the following chapters, one of the main reasons for using an OPF will be to improve the operational efficiency of a distribution network and reduce the power losses within it. In the UK, energy losses in distribution networks account for approximately 6% of the total energy supplied to these networks, which is equivalent to more than 18,000 GWh [106]. Whilst the majority of the energy is lost in the cables that comprise the network, a high percentage is lost in the transformers, especially when they operate close to their capacity.

Also it was reported in a recent study that the majority of the energy losses arise in only a small number of the feeders [107]. According to the study, which analysed a UK network, the two main drivers of the power losses in LV networks were asset utilisation and the length of the feeders. Hence, the reasons for the high energy losses in a small number of feeders can be attributed to the high utilisation of the cables and transformers in those feeders. A more effective distribution of the power could help lower the losses by using more efficiently both the spare capacity in the three phases, as well as in underutilised cables.

When analysing LV networks, there are three levels of latent capacity: 1) the capacity within the phases (whenever one phase is more loaded than the rest); 2) the capacity between feeders supplied by the same substation (one feeder is more loaded than the others) and 3) the capacity between adjacent substations (the transformer of one substation is more loaded). Due to the constant changes in load and the uneven distribution between the phases, these three levels of capacity are constantly changing. Also, different factors affect these capacity levels; the capacity within feeders is impacted by the phase imbalance and the high neutral

current, which increases the thermal constraint in the lines. The substation capacity is affected by the total load supplied, the overall level of imbalance in the transformer and the power losses.

### 2.4.1 Optimal Power Flow for Unbalanced Distribution Networks

As described in section 2.1, the assumption of a balanced three-phase network does not represent accurately the typical behaviour of real distribution networks. The OPF formulation presented in section 2.3 is not well suited to networks with a large phase imbalance. Throughout the literature, there have been several attempts at designing an OPF algorithm that is suitable for three-phase unbalance networks [48, 49, 108, 109]. The study in [48] developed a three-phase OPF with the objective of optimising the control devices present in the network. Rather than formulating an explicit OPF, the algorithm used OpenDSS as a power flow engine and calculated a sensitivity factor of the control variables by evaluating the first and second order derivatives numerically, via small changes. A more comprehensive optimisation framework that includes single-, two-, and three-phase models for the main components of the network and is able to adopt a wider range of objectives was developed by [108]. The study also considers the effect of discrete control devices such as load tap changers (LTC) and capacitors by relaxing the integer variables and transforming the problem from a mixed integer non-linear (MINLP) to a nonlinear programming (NLP) one. The current injection method (CIM), first presented in [110] for single-phase networks and later in [111] for three-phase, follows an approach similar to the one presented in section 2.2 for the power flow. It applies the current injection equations in their rectangular form to each node in the network. Then, the Newton method is used to calculate the nodal currents in the system [111, 112]. The CIM is of particular interest to formulate an OPF, since it allows the full representation of three-phase feeders and both two- and single-phase laterals (the cables connecting a feeder to a customer). Additionally, loads can be modelled as voltage-dependent, and can be connected in either $Y$ or $\Delta$ to their respective phases. The CIM can be formulated by using the sum of the series and shunt injections into the network. That is, the current injected or absorbed by loads/generators connected to a node should be equal to the current flowing into or out of that node to the network through a line. The current mismatch for node $i$ can thus be formulated as the sum of the shunt- and series-connected components as described in (2.33) [109–111]:

$$\Delta I_i = \frac{P_i^{sp} - jQ_i^{sp}}{(E_i)^*} - \sum_{j=1}^{N} Y_{ij} E_j \tag{2.33}$$

where:

$P_i^{sp}, Q_i^{sp}$ = are the specified active and reactive power from a load/generator connected at node $i$

$N$ = is the number of nodes

The voltage at a node $i$ ($E_i$) and the nodal admittance of the line between nodes $i$ and $j$ ($Y_{ij}$) can be expressed in rectangular form as in (2.34) and (2.35):

$$E_i = V_{Re,i} + jV_{Im,i} \tag{2.34}$$

$$Y_{ij} = G_{ij} + jB_{ij} \tag{2.35}$$

The specified active power for node $i$ can be expanded into the generation and load injections, as in (2.36), whereas the specified reactive power has an equivalent representation, as in (2.37).

$$P_i^{sp} = P_{g,i} - P_{d,i} \tag{2.36}$$

$$Q_i^{sp} = Q_{g,i} - Q_{d,i} \tag{2.37}$$

In order to build the three-phase OPF, the current injections are expanded into the real and imaginary parts of its shunt and series components. The real and imaginary parts of the shunt and series components are represented by (2.38) and (2.39), respectively.

$$\Delta I_{Re,i} = \frac{(P_i^{sp})V_{Re,i} + (Q_i^{sp})V_{Im,i}}{(V_{Re,i})^2 + (V_{Im,i})^2} - \sum_{j=1}^{N}(G_{ij}V_{Re,j} - B_{ij}V_{Im,j}) \tag{2.38}$$

$$\Delta I_{Im,i} = \frac{(P_i^{sp})V_{Im,i} - (Q_i^{sp})V_{Re,i}}{(V_{Re,i})^2 + (V_{Im,i})^2} - \sum_{j=1}^{N}(G_{ij}V_{Im,j} + B_{ij}V_{Re,j}) \tag{2.39}$$

Depending on the type of load, the current mismatch can take the three forms for the constant power, constant impedance and constant current loads, as described in equations (2.40), (2.41) and (2.42), respectively [71]. The imaginary part of the current mismatch follows a similar structure and is thus omitted.

$$\Delta I_{Re,i} = \frac{(P_i^{sp})V_{Re,i}^s + (Q_i^{sp})V_{Im,i}}{(V_{Re,i}^s)^2 + (V_{Im,i}^s)^2} - \sum_{j=1}^{N}(G_{ij}V_{Re,j} - B_{ij}V_{Im,j}) \tag{2.40}$$

$$\Delta I_{Re,i} = (P_i^{sp})V_{Re,i} + (Q_i^{sp})V_{Im,i} - \sum_{j=1}^{N}(G_{ij}V_{Re,j} - B_{ij}V_{Im,j}) \tag{2.41}$$

$$\Delta I_{Re,i} = \frac{(P_i^{sp})V_{Re,i} + (Q_i^{sp})V_{Im,i}}{\sqrt{(V_{Re,i})^2 + (V_{Im,i})^2}} - \sum_{j=1}^{N}(G_{ij}V_{Re,j} - B_{ij}V_{Im,j}) \tag{2.42}$$

Given the rectangular representation for the voltage, the upper and lower boundaries need to be expressed as follows:

$$V_{min,i}^2 \leq V_{Re,i}^2 + V_{Im,i}^2 \leq V_{max,i}^2 \tag{2.43}$$

For distribution networks, in general, the voltage at the substation will be balanced, thus the phases will be shifted by 120 degrees [109]. Therefore, the voltage angle constraint is expressed using a similar approach:

$$\theta_{min,i} \leq \tan^{-1}(\frac{V_{Im,i}}{V_{Re,i}}) \leq \theta_{max,i} \tag{2.44}$$

Also, the objective function for the minimisation of the active power losses can be represented using the voltage in rectangular form and the susceptance of the lines. Representing the voltage in its rectangular form has shown to provide several advantages, particularly for IPM methods, one of them being that the objective function has a quadratic form and does not involve trigonometric functions [113]. The sum of all individual line losses is expressed as in (2.45) [98]:

$$\sum_{i=1}^{N}\sum_{j=1}^{N} G_{ij}[(V_{Re,i} - V_{Re,j})^2 + (V_{Im,i} - V_{Im,j})^2] \tag{2.45}$$

Minimising the power losses is chosen as the main objective function, because in general, to reduce the power losses in a given network, other characteristics such as the voltage profile and the phase imbalance need to be improved as well. In other words, by improving the voltage profile along the feeder, reducing the phase imbalance and lowering the loading of the network components, the power losses are minimised as a consequence. For instance, by shifting part of the load from one substation to another, the loading of the transformers is improved, lowering their internal power losses as a result.

The optimisation problem is formulated with the current injection equations for the network as the equality constraints, the operational limits represented as inequalities, and the minimisation of the power losses as the objective, as in (2.46):

$$\underset{x}{\text{minimize}} \quad f(x) = \sum_{i=1}^{N}\sum_{j=1}^{N} G_{ij}[(V_{Re,i} - V_{Re,j})^2 + (V_{Im,i} - V_{Im,j})^2] \tag{2.46a}$$

$$\text{subject to} \quad \Delta I_{Re,i} = 0, \tag{2.46b}$$

$$\Delta I_{Im,i} = 0, \tag{2.46c}$$

$$V_{min}^2 \leq V_{Re,i}^2 + V_{Im,i}^2 \leq V_{max}^2, \tag{2.46d}$$

$$\theta_{min,i} \leq \tan^{-1}(\frac{V_{Im,i}}{V_{Re,i}}) \leq \theta_{max,i} \tag{2.46e}$$

## 2.5 Summary

This chapter introduced the main characteristics of distribution networks and how they differ from transmission networks. It reviewed the key aspects of modelling realistic distribution networks by taking into account characteristics such as the high phase imbalance, the low $X : R$ ratio, the combination of overhead and underground lines and radial operation. Also, this chapter introduced the typical power flow and OPF formulations and highlighted the ways in which the differences between transmission and distribution networks influence the assumptions that can made when designing these algorithms. The OpenDSS software tool was also introduced as an implementation of the power flow algorithm that can be used to model realistic distribution networks. The fixed-point algorithm used by OpenDSS for the solution of the power flow was described, as it represents a computationally efficient solution method, which will be leveraged in subsequent chapters. Finally, with an OPF framework established which can be applied to distribution networks, the next chapter will describe the changes needed in the OPF formulation to accommodate new components, such as power electronics devices and DC connections.



# Chapter 3

# Active Network Management

Some of the results in this chapter are taken from the following papers:

- J. Perez-Olvera, T. C. Green, and A. Junyent-Ferre, "Using multi-terminal DC networks to improve thehosting capacity of distribution networks," in *2018 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, Conference Proceedings, pp. 1–6
- J. Perez-Olvera, T. C. Green, and A. Junyent-Ferre, "Active network management in LV networks: a case study in the UK," in *Power & Energy Society General Meeting (PESGM)*, 2020, pp. 1–5.

## 3.1 Active Network Management

Active Network Management (ANM) is one of the tools that can be used by DNOs to deal with the large-scale integration of LCTs in distribution networks and continuously manage the network capacity.

The Energy Networks Association (ENA) in the UK defines ANM as [114]:

> *Using flexible network customers autonomously and in real-time to increase the utilisation of network assets without breaching operational limits, thereby reducing the need for reinforcement, speeding up connections and reducing costs. The definition can be extended to include DNO-owned network assets such as transformers and Power Electronic Devices to provide a wider range of network services.*

Despite leaning towards the management of customers flexibility, the ENA's definition of ANM also contemplates the use of network assets, such as power electronic devices, to provide network services, as part of an ANM scheme.

Previous research on ANM has focused on controlling the output of solar PV generators [115] and other forms of DG [116, 117] to relieve network constraints. In [115], the authors proposed a distributed control

to manage the active and reactive output of solar PV generators at the inverter level. The scheme managed shows an acceptable performance with minimal communication requirements. However, the study does not include a realistic model of the LV network that accounts for the intrinsic phase imbalance of these networks. Instead it assumes the PV generators inject the same amount of active and reactive power in all phases. As was discussed in section 2.1, the variation in voltage between phases can be large and injections in one phase can affect the voltage in the other phases. A field trial of an ANM system was carried out in [117] to control both a solar PV array and a wind generation unit. The goal of the ANM scheme was to monitor the network and manage the power output of both generation units to avoid exceeding the upper voltage constraints. The voltage at the key constraint points were constantly monitored and compared to a predetermined threshold. The generator output was then adjusted in discrete steps. It is worth noting that despite the benefits of managing the output of generation units, these schemes cannot increase the overall asset utilisation between phases or between feeders. Instead they only make a better use of the generated energy across time to limit curtailment or to avoid network constraints. This is due to the fact that power cannot be redistributed between different assets (i.e. cables, transformers, etc.) unless a connection exists. For instance, the ANM schemes in [117] and [115] did not contemplate the possibility of meshing the neighbouring substations in order to export the excess power from the DG to supply the loads in a different substation, rather than capping the DG output.

On the other hand, ANM can also include the control of the power flows within a network, or a part of it, by controlling the DNO assets directly [114]. As discussed in section 1.4, a network can be meshed to decrease the impact of DG and LTCs in general. By creating alternative paths for the power flows, the loading can be shared between substations and feeders. However, as it was noted, the power flow is governed by the voltage along the feeders and the network impedance. A highly loaded substation can become even more loaded as a result of the meshing if key factors are not taken into account. Understanding the circuit layout, fault levels and location of the loads is essential to predict the power transfers accurately.

A continuous control of the power flows can instead be achieved through the interconnection of feeders and substations using power converters. Given the $X : R$ ratio of distribution networks, managing the active power can have a greater effect on the voltage than compensating with reactive power, as it is normally done in transmission networks. However, unlike managing the output of a generator or the demand from a customer, when actively managing the power flows within a network it is essential to model the cables and transformers in detail, to observe the effect of the converters on the overall network operation. As described in section 2.2, for distribution networks, this means having accurate network models that reflect their main characteristics: high resistivity, uneven load distribution, high phase imbalance and radial operation. Thus to design a control for an ANM scheme, which is capable of managing the power flows, a detailed model of the network and of the power converters is imperative.

One of the main challenges in developing active control of distribution networks is the lack of instrumentation and monitoring, which affects the visibility, as described in section 1.3 This lack of information about the loading state of the individual components of the network makes it difficult to implement any type of centralised control for ANM. The use of smart meters data alone cannot provide the level of detail required to implement an ANM solution, such as power flow management [26]. Nonetheless, with the rapid uptake of LCTs as highlighted in section 1.2, improving the visibility of these networks has become a necessity for DNOs.

In recent years, several DNOs in the UK have shown interest in installing more monitoring devices in LV networks to improve their observability and to underpin the deployment of smart grid solutions. For instance, the project in [118] trialled the installation of monitoring devices at secondary substations and at some feeder end-points. The data was used to categorise the customers within the network, analyse load aggregation patterns at substation level and build a forecasting model for domestic consumption based on the data collected. To analyse the accuracy of the assumptions made about the impacts of PV connections to the LV network, the project [22] monitored several secondary substations, as well as PV installations connected downstream from these substations, for more than 12 months. One of the main issues found by the project was that there was a significant level of uncertainty on the current level of PV connections, the phase to which they are connected and their size, compared to the DNO database/registrar. On the other hand, to establish the business case for improved visibility in LV networks, the OpenLV project developed a platform for the deployment of software solutions in the form of *apps* [119]. To achieve this, the host DNO installed monitoring devices at secondary substations and heads of feeders, which would give an accurate estimate of the loading state of secondary substations, as well as highlight any potential balancing opportunity. Additionally, a network meshing approach using controllable circuit breakers was trialled to release capacity between substations.

Two LV substations were connected via one of their outgoing feeders using a normally open point (NOP) and ALVIN Reclose™ devices (a type of controllable circuit breaker [55]). Monitoring was deployed at the LV side of the transformer and at the head of each of the outgoing feeders. Fig. 3.1 shows an overview of the monitoring available and the ANM solution implemented as part of the OpenLV project [55]. Remote terminal units (RTUs) were used to monitor and collect granular time-series data such as: voltage, current and power [55, 120]. This information proved useful in determining which substation has the most loading and at what time and if some other substation has spare capacity. Despite the benefits of this type of ANM, the long-term advantages are not clear, as there is a lack of control over the power flows.

Given the evidence presented in 2.1, where it was shown that the LV network is almost 100 times in equivalent length than the HV network [37], together with the evidence in this section, it is unlikely that full observability of LV networks will be attainable in the near future. Mainly due to their fragmentation and the

high cost. However, as the level of LCTs increases, putting pressure on DNOs, it is likely that monitoring at key points will be available in the majority of the LV networks, as illustrated in Fig. 3.1. This will support the adoption of ANM solutions based on power converters by enabling the development and implementation of sophisticated control strategies for real-time operation.

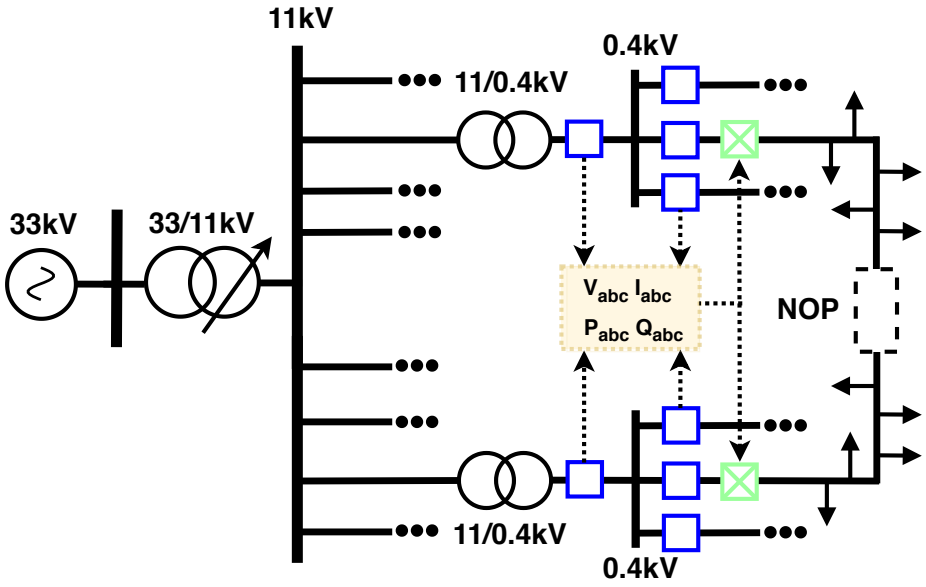



Figure 3.1: Overview of network meshing in the OpenLV project. The monitoring units are shown in blue and the ALVIN Reclose™ devices used for ANM are in green [55].

## 3.2 Active Network Management Using Power Electronics Devices

As introduced in section 1.4, power electronic converters can be used to control the power flows in a distribution network, as part of an ANM strategy, and release latent capacity at different levels (e.g. between phases, feeders and substations), as described in section 2.3. Previous research has explored the potential benefits of including stand-alone power converters in DNs [121]. The study in [122] analysed how several solutions based on power converters can delay the occurrence of problems with thermal and voltage limits in DNs with positive results. Of the various power converters that can be used, the most versatile and simple device that has gained popularity in recent years is the voltage source converter (VSC), in a back-to-back configuration comprising two or more converters. This configuration has been referred to using several names throughout the literature: Intelligent nodes [123], loop-balance controllers [124], DC-links [125], Soft-Open-Points (SOPs) [126], etc. However, the basic structure remains the same, a number of VSCs that serve to connect together parts of a medium-voltage (MV) or low-voltage (LV) network and that offer complete control over the active and reactive power to and from the converters. Fig. 3.2 shows the two LV substations from the previous section connected using power converters instead of meshing via their NOP. In this way, the power

injections can be controlled to enhance the operation of the network.

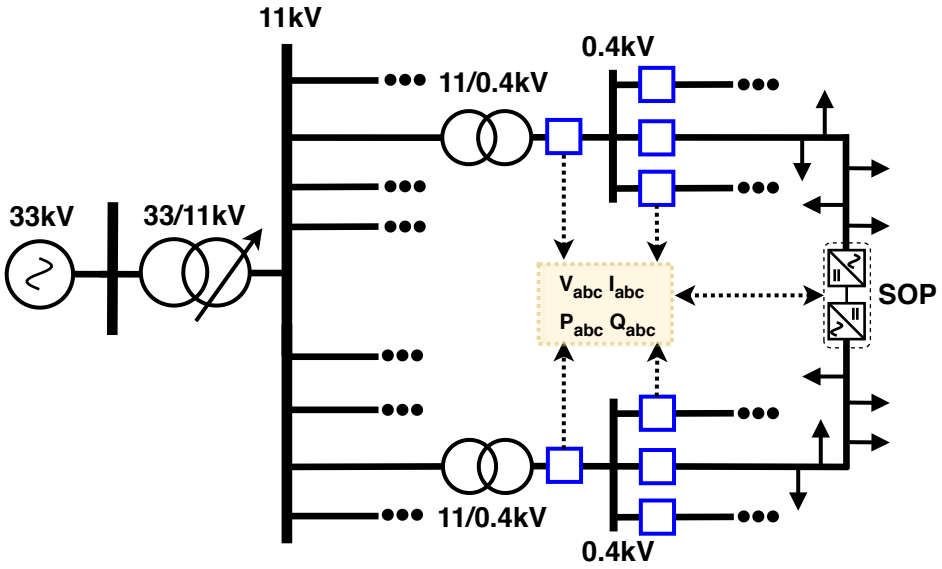



Figure 3.2: Two LV substations connected using power electronic converters, with the monitoring units shown in blue. This back-to-back connection between the converters is also referred to as SOP.

In one of the earliest studies [123], the authors proposed meshing two distribution feeders at their end points via individual back-to-back converters forming a common DC-link, enabling a full control of the power flow between the feeders. The main goal of the proposed scheme was to regulate the voltage of the feeders and to ensure it remained within the specified limits. One of the main limitations of this study is that the power converters were modelled separately and their optimal settings were obtained sequentially for one converter at a time rather than calculating all of them at the same time. Also, the power losses of the converters and realistic load profiles were not considered. The study [125], proposed an optimisation algorithm to assess the technical and economic benefits of using DC links to mesh feeders in distribution networks. In a follow-up study [127], the authors focused on analysing the advantages brought by the introduction of DC-links to the real-time operation of the network. The study modelled the performance of the link in both steady-state and transient operation, when it was used primarily to minimise the power losses in the network. Despite the positive results obtained in both the benchmark network and the test system used, the study neglected the typical phase imbalance characteristic of distribution networks, as described in section 2.2. The phase imbalance adds another layer of complexity to the design of a suitable control system and needs to be considered. The results of the study in [122] showed that power converters can be used effectively to reduce, and in some cases eliminate, over-voltage problems. It was noted by the authors that the benefits vary from one network to another and that placing the power converters in the right location is key to maximise the benefits. The work in [128] proposed using DC links to mesh two real DNs, situated in an urban location. In this case, the authors employed a realistic model of the networks, built in OpenDSS. The results showed that

the link successfully equalised the overall load across the feeders. The increased reliability and power quality are two of the main benefits highlighted from the incorporation of power converters. Also, the results suggested that there is a significant opportunity in meshing networks with non-coincidental load profiles. Nonetheless, the authors did not represent the converters losses nor the DC connections. Moreover, only power flow calculations were performed, which can only output the feasible set points for the converters but not the optimal.

Similarly, a field trial of the SOP concept was recently carried out by a DNO in two cities in the UK. The converters were used to connect two feeders supplied from different substations, improving the operating characteristics of the network [129]. The control proposed, based on a hysteresis method, used both local and remote measurements to determine the power set-points for the converters. However, the objective of the control did not consider global network conditions such as the power losses or the phase imbalance. This is essential when aiming to improve the overall utilisation of the network components, especially in LV networks. The work by [130] addressed the challenge of finding optimal locations and sizing DC-links for distribution networks. The findings of the study challenge the notion that the best location for DC-links is in the place where NOPs are usually located. From the results, it can be observed that depending on the network, the optimal topology that would lead to the lowest total cost could be different that initially thought. The study also assumed that the DC-links can only mesh one pair of feeders together, neglecting the possibility of having multi-port DC-links or multi-terminal DC (MTDC) grids that link more than two feeders and share a common DC bus, an approach that could potentially be more effective and help reduce the number of VSC converters needed. Also, individual converters could be added incrementally to the most problematic feeders and used to mesh them together to an existing DC bus. This way, the power losses arising from the converters would increase only slightly whilst the overall power losses would be expected to decrease substantially. The authors in [131] proposed connecting the feeders of a network using power converters and calculating the set points via an OPF. However, the study only considered point-to-point connections, rather than a MTDC. Moreover, only the single-phase equivalent of the network was used. In [132], the study showed that an improved network operation was achieved by the MTDC network, which connected all the end points to each other, instead of only in pairs, as originally proposed in [131]. The difference was especially notorious on the overall power losses with a further decrease of approximately 10%.

Another possibility worth considering is that once two or more DC-links exist within the network, they can be connected to each other via a DC line. These two factors could strongly benefit the operation and regulation of the network if implemented successfully. Although, the concept of a low-voltage DC (LVDC) network has been explored before [133], it has mostly been treated as an alternative that would replace in its entirety the current low-voltage networks. However, an incremental deployment of small DC grids, derived from already-present LV feeders might be more realistic. Having a LVDC network could unlock several

advantages when considering the combined effects of embedded LCTs, especially at high penetration levels. One advantage would be the increased efficiency by the reduction in the number of conversion stages needed to connect DC-based LCTs. The work in [134] compared the efficiency of a proposed DC network composed of both medium- and low-voltage levels against a hybrid AC/DC network and a conventional AC configuration. The results showed that the conversion of part of the network to DC, especially the low-voltage feeders where the loads are located, seemed to improve significantly the utilisation of the main transformer, releasing extra capacity due to the absence of reactive power flows. In other words, the capacity of the network to host more loads, and potentially even more DG, could be enhanced without overloading the main transformer.

Overall, these studies provide evidence of the effectiveness of incorporating Voltage Source Converters (VSCs) in DNs to improve the operation of these networks. However, the advantages of extending this concept to a more complex network such as a multi-terminal DC (MTDC) network, one that is formed of several VSCs, has not been investigated thoroughly. Research is needed to compare the advantages of multi-port DC links against the conventional dual-port approach and the benefits of having a common DC bus versus separate DC-links. Moreover, the design of a suitable control for real-time operation that is able to produce optimal set-points for the converters on a continuous basis has not been addressed in the literature. The problem of scheduling the power set points for multiple converters, whilst considering the physical constraints of the distribution network and the load/generation present can be seen as a nonlinear optimisation over time and thus can be formulated as an extension of the OPF algorithm presented in Chapter 2.

## 3.3 Active Network Management Control Using AC/DC OPF

In order to analyse the operation of hybrid AC/DC networks, which have embedded power converters and DC connections, a power flow algorithm must take into account both the AC part and the DC part of the network. The modelling framework introduced in sections 2.2 and 2.3 can be used as the starting point for such an OPF formulation suitable for realistic distribution networks with embedded power converters and DC connections. Such a formulation is referred to as AC/DC OPF, as it models both the AC network and the multiple DC links formed by connecting the power converters together.

There have been several attempts in the literature to formulate a power flow algorithm that would deal with a MTDC grid connected to an AC network, whilst including a detailed model of the VSCs [135, 136]. The study in [136] used a sequential approach, in which the power flows of the DC and the AC network are solved independently but the outputs of DC power flow are used as the inputs for the AC power flow, and vice versa, such that each updates the other at each iteration until a converged solution is found. In a subsequent study, the authors [137] adapted the power flow algorithm to incorporate the option of having

a distributed DC voltage control for the VSCs. By including a droop characteristic as part of the control scheme of the VSCs, the task of regulating the DC voltage was shared among the converters forming the MTDC network, as opposed to having a centralised configuration where one converter acted as the reference. Although a power flow is useful for determining the state of the network at a certain point in time, in order to calculate the ideal operating point (i.e. the one with the minimum power losses), and to calculate suitable set-points for the converters, an OPF formulation is required. To incorporate the VSCs, the standard AC OPF needs to be altered to include the equations of both the converters and the power flows in the DC network. For instance, the power balance for each VSC needs to be included as an equality constraint. Recent studies have proposed a mixed AC/DC OPF formulation to analyse MTDC networks connecting offshore wind farms to the main AC transmission system through HVDC systems [138, 139]. Following a similar approach, the power flow balance for each of the converters in the network can be expressed as in (3.1).

$$P_{injVSC,i} - P_{lossVSC,i} - P_{DC,i} = 0 \tag{3.1}$$

where:

$P_{injVSC,i}$ = is the power going into the AC side of the converter $i$

$P_{lossVSC,i}$ = is the power losses of the converter $i$

$P_{DC,i}$ = is the power going out of the DC side of converter $i$

Fig. 3.3. illustrates the different variables that need to be incorporated into the power flow equations.

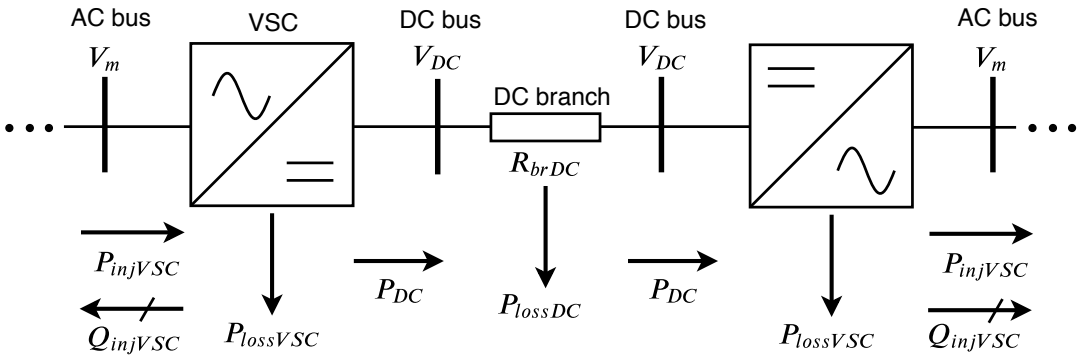



Figure 3.3: VSCs connected back-to-back through their respective DC bus forming a SOP.

The losses of the converters are normally expressed as a polynomial that depends on the current as in (3.2) [140]:

$$P_{lossVSC,i} = a + bI_{VSC,i} + cI^2_{VSC,i} \tag{3.2}$$

Here, the three coefficients ($a$, $b$, $c$) in (3.2) represent the no-load, switching and conduction losses of the converters, respectively [140]. The no-load losses are usually those associated with auxiliary equipment required by the converters to operate (e.g. cooling equipment). The switching losses are those related to the

switching of the individual valves [138, 140]. At rated power, the losses for a two-port 240 kVA converter were reported to be around 7% (16.8kW) of the injected power, whereas at 20% of its rated power (48 kW) the losses represent approximately 15% of the injected power (7.3 kW) [62].The current through the converter is expressed as a function of the active and reactive power injections and the voltage at the AC node where the converter is connected, as in (3.3). Thus, the losses of the converters vary as the power injections change.

$$I_{VSC,i} = \frac{\sqrt{P_{injVSC,i}^2 + Q_{injVSC,i}^2}}{V_{AC,i}} \tag{3.3}$$

where:

$Q_{injVSC,i}$ = is the reactive power injection into the AC side of the converter $i$
$V_{AC,i}$ = is the voltage at the AC side of converter $i$

The power injections inside the MTDC network can be expressed in a similar way as those in the AC network, by using the admittance matrix [136]. One of the converters forming the MTDC acts as a slack bus keeping the voltage at 1.0 pu. Fig. 3.4. shows an example of the flows inside a generic MTDC network and (3.4) summarises the flows for all branches.

$$P_{DC,i} = V_{DC,i} \sum_{j=1}^{N} V_{DC,j} G_{DC,ij} \tag{3.4}$$

where:

$N$ = is the number of nodes in the DC network
$V_{DC,i}$ = is the voltage at the DC side of converter $i$
$G_{DC,ij}$ = is conductance of a DC line between nodes $i$ and $j$

In the AC/DC OPF the power injections from the VSCs are added to the set of control variables that are used to achieve the optimisation objective [141]. Given the characteristics of the VSCs, their contribution can be modelled as a shunt current injection into a particular node, similar to a load/generator and thus can be incorporated into the specified active $P_i^{sp}$ and reactive power $Q_i^{sp}$, as in (3.5) and (3.6).

$$P_i^{sp} = P_{g,i} - P_{d,i} - P_{injVSC,i} \tag{3.5}$$

$$Q_i^{sp} = Q_{g,i} - Q_{d,i} - Q_{injVSC,i} \tag{3.6}$$

The rating of the converters can be expressed as in (3.7), and incorporated into the AC/DC OPF.

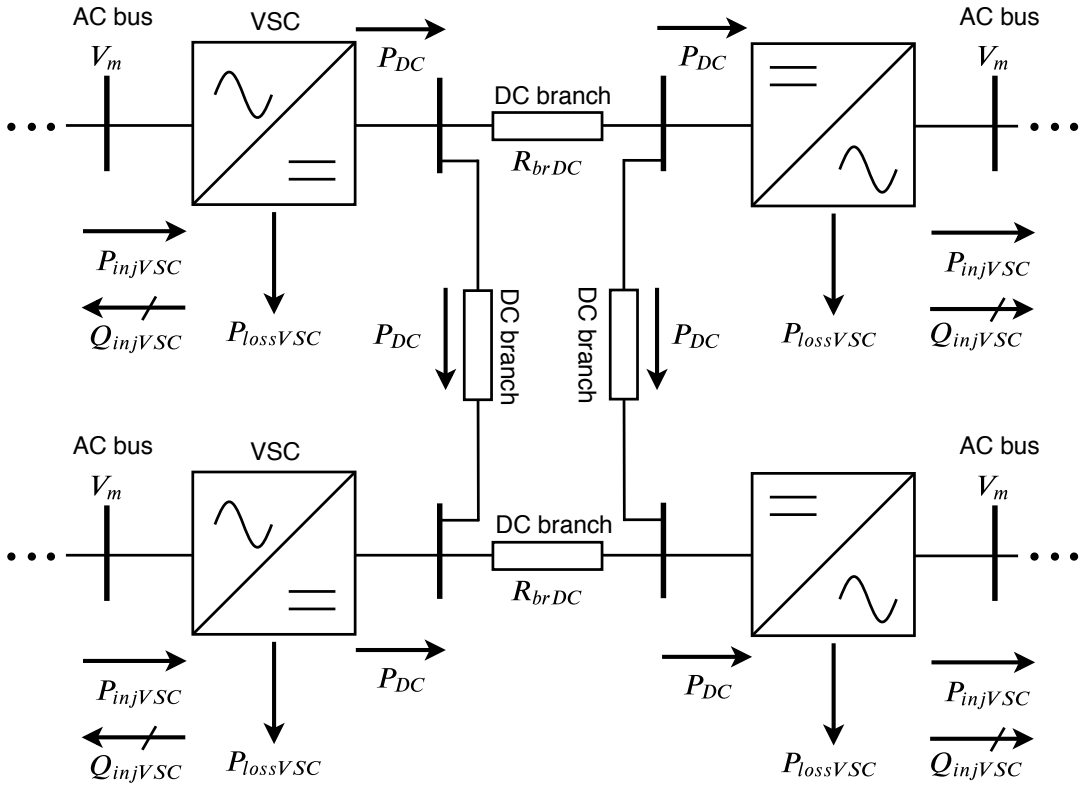



Figure 3.4: Multiple VSCs forming a MTDC network.

$$S_{VSC,i} = \sqrt{P_{injVSC,i}^2 + Q_{injVSC,i}^2} \tag{3.7}$$

Additionally, the power flowing in the MTDC network needs to be included as a separate equality constraint together with the converter losses and the power balance for each of the converters. The equations introduced in the previous section for the converters are used as equality constraints for the AC/DC OPF. Hence, the AC/DC OPF gives as a result the optimal active and reactive power set-points for the VSCs.

Unlike the formulation used for HVDC transmission networks, where one of the main assumptions made is that of a balanced operation, in this case the high phase imbalance is considered and all three phases are modelled explicitly. For the purpose of this case study, the minimisation of the power losses will be considered as the main objective function. As described in section 2.5, by minimising the power losses, other parameters such as, the voltage profile, the phase imbalance and the overall loading of the network components are also improved. This objective function can be expressed as in (2.45), or as the difference between the power supplied by the substation and the power injections of the generators and those of the loads at all the nodes $N$, as in (3.8).

$$f(x) = \sum_{i=1}^{N} (P_{g,i} - P_{d,i}) \tag{3.8}$$

Thus, the AC/DC OPF can be formulated as in (3.9):

$$\underset{x}{\text{minimize}} \quad f(x) = \sum_{i=1}^{N}\sum_{j=1}^{N} G_{ij}[(V_{Re,i} - V_{Re,j})^2 + (V_{Im,i} - V_{Im,j})^2] \tag{3.9a}$$

$$\text{subject to} \quad \Delta I_{Re,i} = 0, \tag{3.9b}$$

$$\Delta I_{Im,i} = 0, \tag{3.9c}$$

$$P_{injVSC,i} - P_{lossVSC,i} - P_{DC,i} = 0 = 0, \tag{3.9d}$$

$$V_{min}^2 \leq V_{Re,i}^2 + V_{Im,i}^2 \leq V_{max}^2, \tag{3.9e}$$

$$\theta_{min,i} \leq \tan^{-1}(\frac{V_{Im,i}}{V_{Re,i}}) \leq \theta_{max,i} \tag{3.9f}$$

where the real and imaginary parts of the shunt and series injections, for constant power loads, are formulated as (3.10) and (3.11), respectively, as described in section 2.3. The converters injections are included in the specified active $P_i^{sp}$ and reactive power $Q_i^{sp}$, as in (3.5) and (3.6).

$$\Delta I_{Re,i} = \frac{(P_i^{sp})V_{Re,i} + (Q_i^{sp})V_{Im,i}}{(V_{Re,i})^2 + (V_{Im,i})^2} - \sum_{j=1}^{N}(G_{ij}V_{Re,j} - B_{ij}V_{Im,j}) \tag{3.10}$$

$$\Delta I_{Im,i} = \frac{(P_i^{sp})V_{Im,i} - (Q_i^{sp})V_{Re,i}}{(V_{Re,i})^2 + (V_{Im,i})^2} - \sum_{j=1}^{N}(G_{ij}V_{Im,j} + B_{ij}V_{Re,j}) \tag{3.11}$$

## 3.4 Case Studies

This section will demonstrate the performance of the AC/DC OPF algorithm through case studies. The AC/DC OPF algorithm was implemented in MATLAB and solved with the *fmincon* function, using the built-in IPM solver. The IPM solver calculates the gradients of the constraints via finite differences, a type of numeric differentiation [142]. The results obtained from the case studies were first validated with those obtained using OpenDSS without embedded converters [95]. First, the operational benefits of embedded power converters will be demonstrated, particularly when a large volume of LCTs is present in the network. Then, the OPF suitability as a control for ANM in real operation will be discussed, with emphasis on it is accuracy, its computational tractability and its observability requirements.

The different case studies will be used to represent: a) distribution networks with both underground cables and overhead lines, as is the case of the IEEE 123-bus test feeder, b) the impact of the unavailability of real monitoring data at low-voltage level, and c) the topology of LV networks in Europe. The two last case studies also highlight the challenge of modelling LV networks, in the case of the Lotherton and Canefields

substations, real data is available for the aggregated loads, however the topology and connectivity of the network is not available. On the other hand, the topology of the European LV network is available but the individual load profiles are not.

### 3.4.1 IEEE 123-bus Test Feeder

To validate the performance of the AC/DC OPF formulation, a benchmark distribution network from the IEEE Distribution System Analysis Subcommittee Working group is used first. This radial IEEE test feeder with full three-phase characteristics consists of several branches with different lengths and is supplied by a main substation with a 5 MVA 115/4.16 kV transformer [143]. The operating voltage of the feeder is 4.16 kV and the peak load is 3.49 MW and 1.92 Mvar comprising 114 individual loads. All three types of voltage-dependent load are represented in this feeder. Also, both overhead and underground lines with different phasing are present. For the purpose of this analysis, both the shunt capacitor banks and the voltage regulators were omitted. The reason for this was to observe the enhancement of the operating conditions when a the power converters were included. The diagram of the test feeder is shown in Fig 3.5 [143, 144]. A power flow is run with the default loads for each hour time-step in order to obtain the base-case conditions for the network. The voltage unbalance for this case is shown in Fig. 3.7.

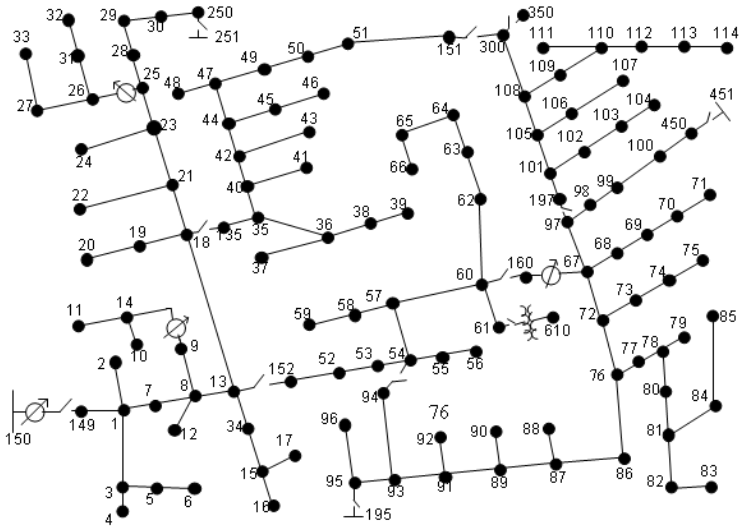



Figure 3.5: IEEE 123-bus test feeder diagram [143, 144].

PV generators with a capacity of between 15 kW and 120 kW, were assigned randomly to 35 nodes in the network. Some of them are connected to a all the three-phases and some only to a single phase. It will be assumed that they operate at a power factor of 1. Additionally, groups of 10 EVs, each group with a peak demand of 33 kW, will be connected at various different nodes in the network. Typical hourly profiles for the loads, EVs and the PV generators are used from two projects in the UK [22, 145] and are shown in Fig. 3.6. By comparing Fig. 3.7 and Fig. 3.8 it can be seen that the inclusion of PV generators and EVs results in a

significant increase in the voltage unbalance between the phases, especially at peak generation and peak load hours.

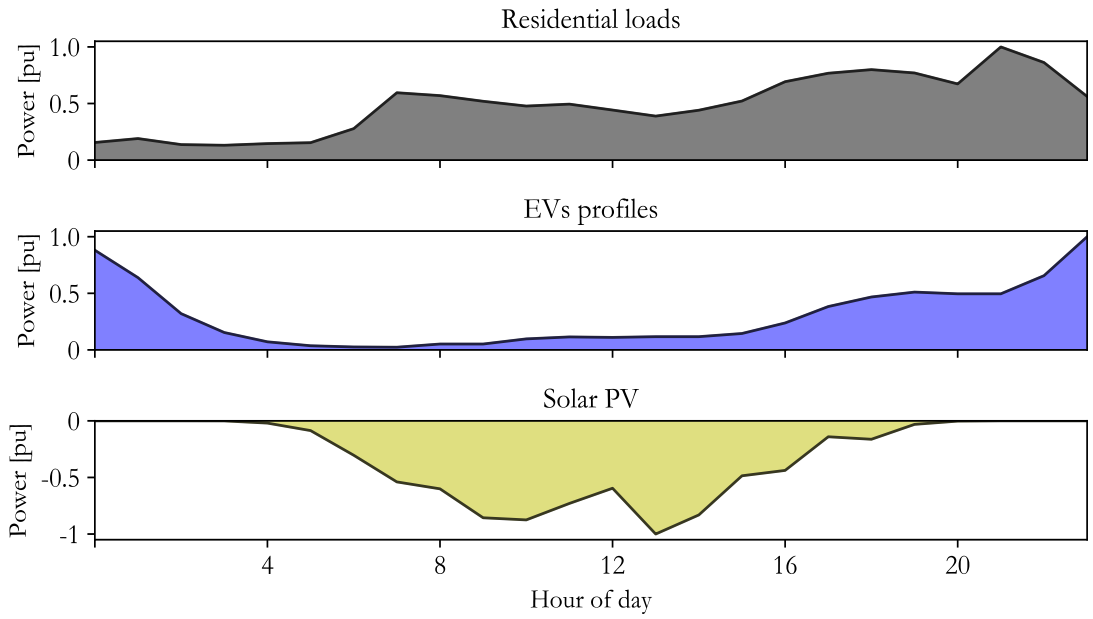



Figure 3.6: Load and generation profiles with 1-hour resolution.

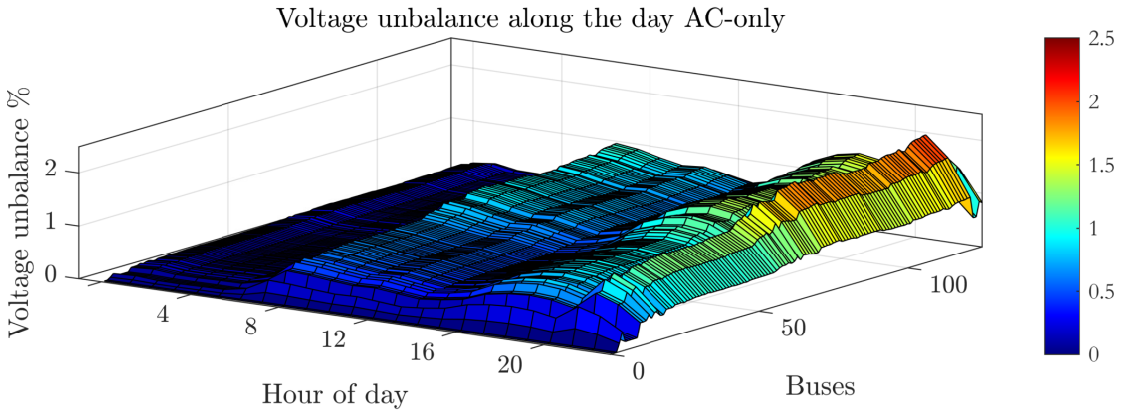



Figure 3.7: Voltage unbalance with default loads only.

Fig. 3.9 depicts the current supplied by the substation. It can be seen that there is a large difference between the currents of each phase, mainly during the hours of peak solar PV generation (08:00-16:00 hours). This is because a large fraction of the solar PV is connected to phase C and with low loading conditions it results in a significant reverse power flow. The large difference of current between phases results in a large current in the neutral, which is observed in Fig. 3.9. The significant reverse power flow in phase C, during the hours of peak solar PV generation, is observed in Fig. 3.10.

A breakdown of the active and reactive power losses by component type is shown in Fig. 3.11. In this case, the reactive power losses are higher compared to the active power losses. Also, the losses in the lines are higher than those in the transformer. The total losses of the lines and transformer for the day are 1943.89 kWh.

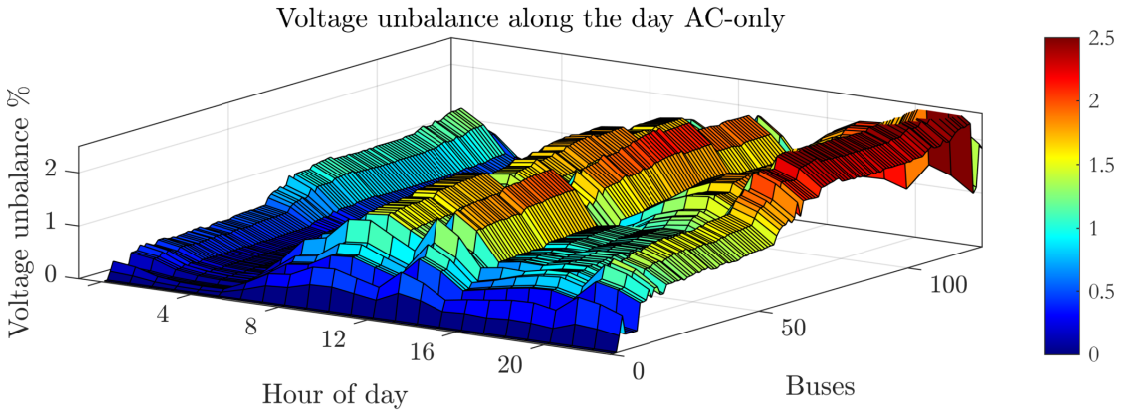



Figure 3.8: Voltage unbalance with default loads, EVs and solar PV generators.

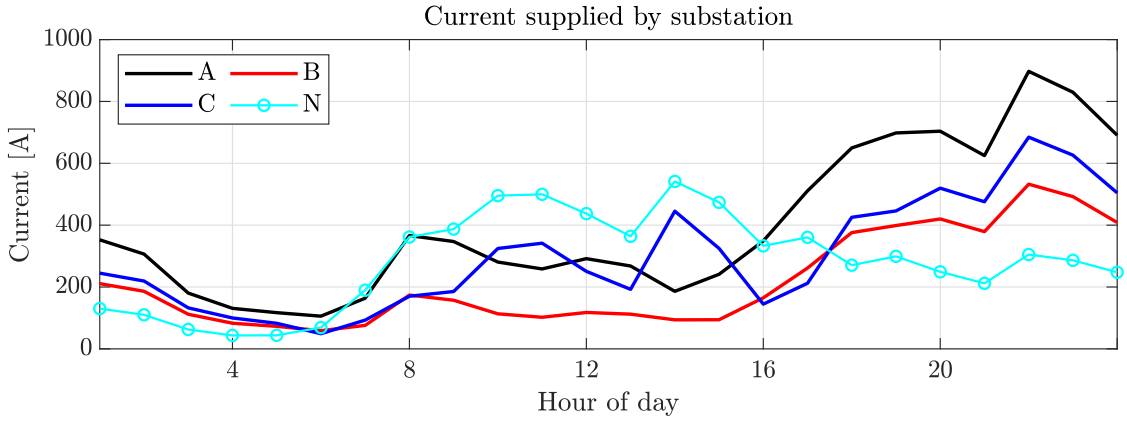



Figure 3.9: Current supplied by substation with default loads, EVs and solar PV generators.

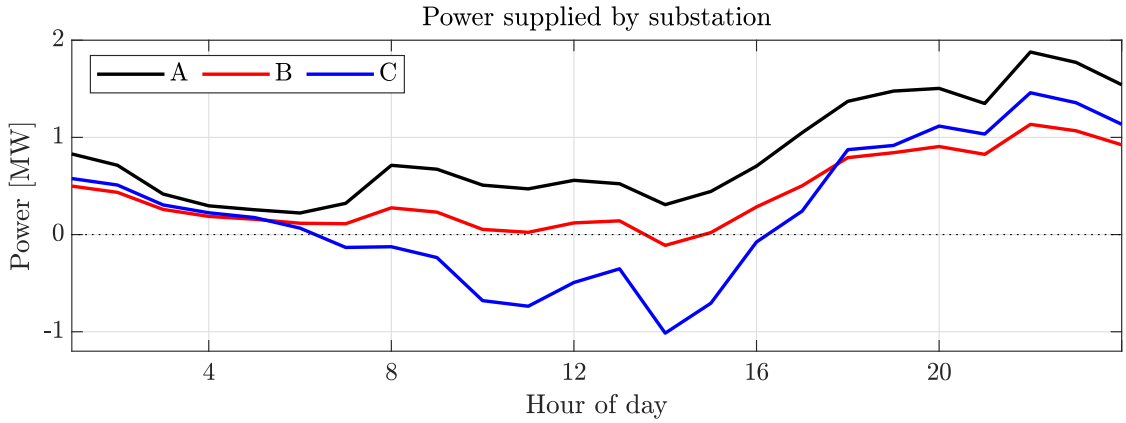



Figure 3.10: Power supplied by substation with default loads, EVs and solar PV generators.

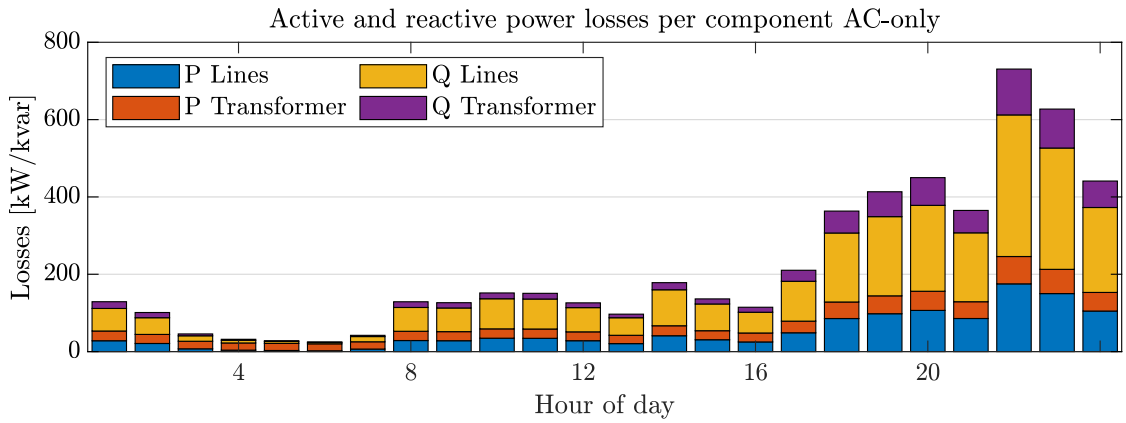



Figure 3.11: Active and reactive power losses by component with default loads, EVs and solar PV generators.

It is worth noting that computation of a power flow in OpenDSS takes a few milliseconds for each time-step, as seen in Fig. 3.12. On average, a power flow takes 47 ms to run, denoted by the red dashed line in Fig. 3.12 for a total run-time of 1.025 s. Given the nonlinear characteristics of the loads, as described in section 2.2, a different number of iterations of the fixed-point algorithm might be required for each time-step to converge.

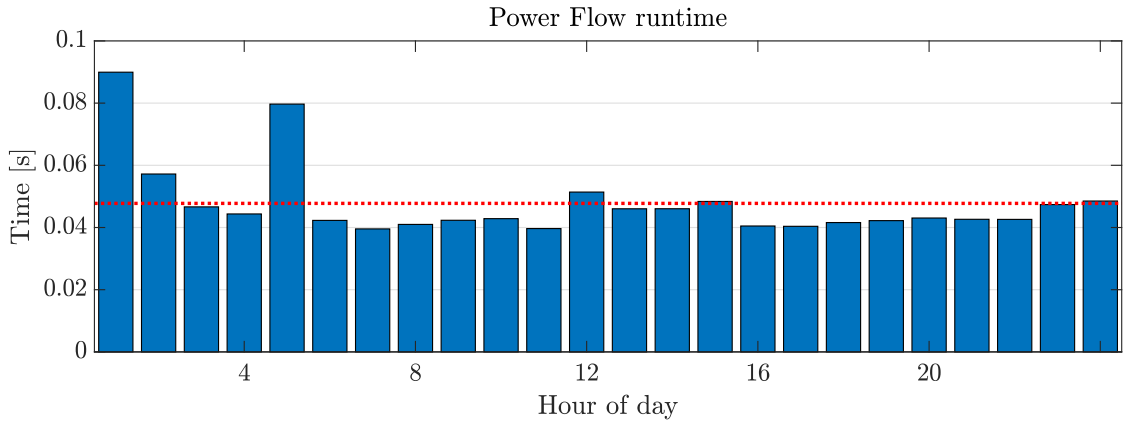



Figure 3.12: Runtime of the power flow for each time-step, the dashed line represents the average run-time.

To explore the potential improvement in the operation of the test feeder, a MTDC network is proposed, which can connect node 151 to node 300, connect node 39 to node 66, and also connect these two pairs of nodes. The location of the MTDC network was chosen based on the close physical proximity of the nodes in question and the presence of an existing normally open point between nodes 151 and 300. These factors would make it viable in practice to install the pairs of converters. These nodes are depicted in red in Fig. 3.13 and a detailed illustration of the internal and external connections between converters is shown in Fig. 3.14. The converters single phase rating is 250 kVA with a power factor of 0.6, which can be either lagging or leading depending on the needs of the network. This power factor was selected given the $X : R$ characteristic of the network and the higher reactive power losses observed in the base case. The AC/DC OPF will be run for every time step of the hourly load profile to obtain the converter's set points, using the results of the power flow as initial point for the optimisation. It has been shown that the initial point chosen

can greatly influence the behaviour of the IPM algorithm and that if possible this initial point should be as close as possible to the optimal [105]. The network parameters (i.e. the voltage and current at all nodes) calculated by the power flow should represent a good approximation and has been used in previous studies with good results [100, 113]. It is worth noting that obtaining the optimal set points of the converters is not trivial and if the power injections are smaller or larger or not adequately coordinated, they could end up worsening the operation of the network. Also, the optimisation minimises the overall power losses whilst taking into account the inherent power losses of the converters thus adding an extra layer of complexity.

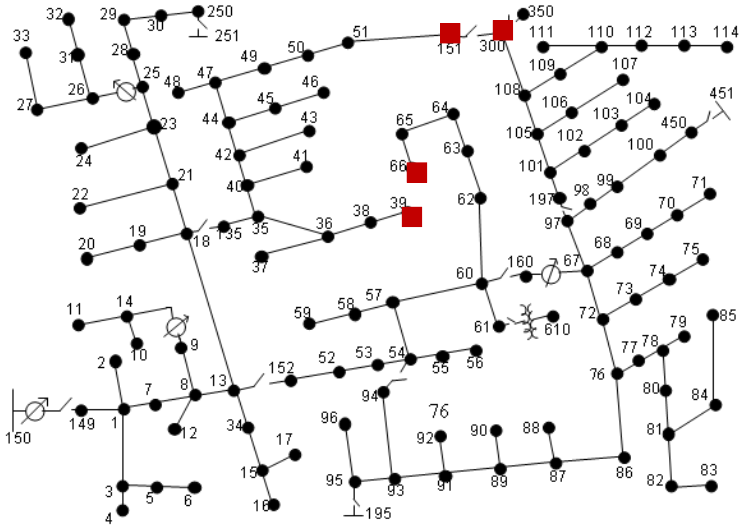


Figure 3.13: IEEE 123-bus test feeder diagram with VSCs at marked buses (39, 66, 151, 300) [143, 144].

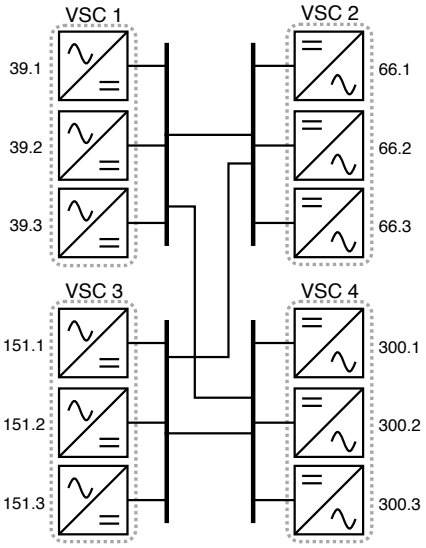



Figure 3.14: IEEE 123 MTDC diagram. Each three-phase converter is composed of three single phase converters that share a common DC bus.

When comparing Fig. 3.8 and Fig. 3.15, it is clear that the converters help distribute the power more efficiently, thus decreasing the voltage unbalance across the network. Especially during the hours of peak load, without the converters the voltage unbalance is higher than the recommended 2%, reaching 2.3% in

some of the buses. However, when the converters are active, the voltage unbalance is maintained below 1% for most of the day, with only a small peak of 0.8% around the hours of maximum PV output and peak load.

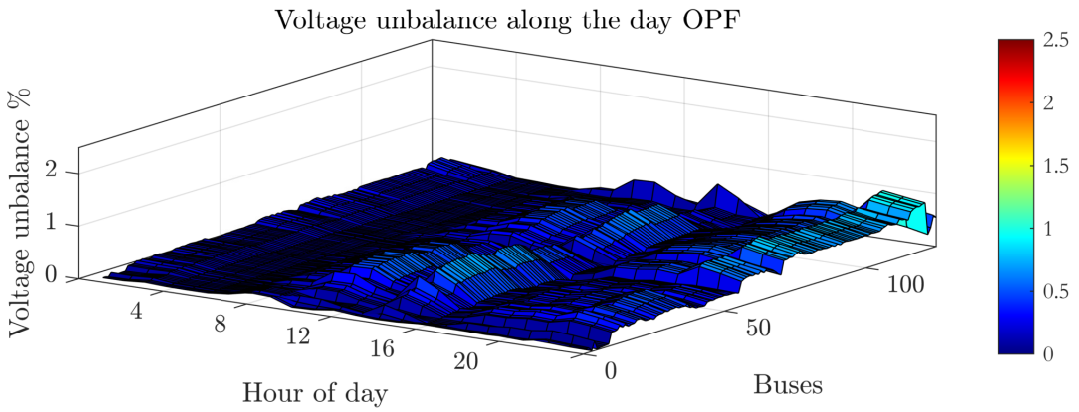



Figure 3.15: Voltage unbalance with default loads, EVs and solar PV generators with AC/DC OPF.

Also, this is observed by comparing the current supplied by the substation in Fig. 3.9 and Fig. 3.16. The current in the neutral is reduced significantly when the converters are operating, from an average of 282 A to only 76 A, a reduction of more than 70%. Both the amount and the phase difference of the reverse power flow are reduced significantly from Fig. 3.10 to Fig. 3.17, which contributes to the overall release of latent capacity.

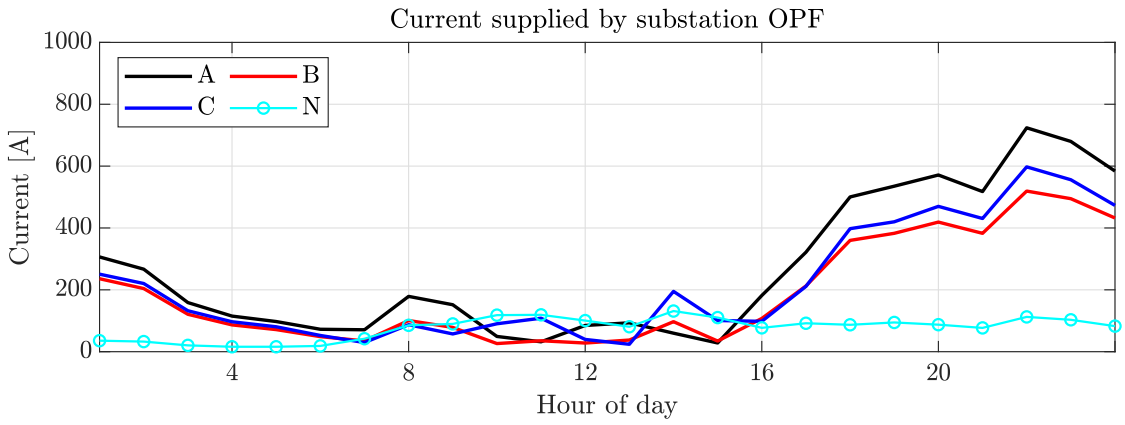



Figure 3.16: Current supplied by substation with default loads, EVs and solar PV generators, calculated by the OPF.

The AC/DC OPF determines the active and reactive power set points for each of the converters and also it calculates how the power should be distributed between the converters and between the phases, as seen in Fig. 3.14. The active and reactive power injections from the power converters are illustrated in Fig. 3.18 and Fig. 3.19. These two figures show the total power going into, or out of, each of the individual phases. It is clear how some of the converters shift their power injections when the output of the solar PV generation is at its peak. As seen in Fig. 3.19, the converters inject reactive power throughout the day, following the overall load profile of the network, to compensate for the reactive power needs of the transformer and lines.

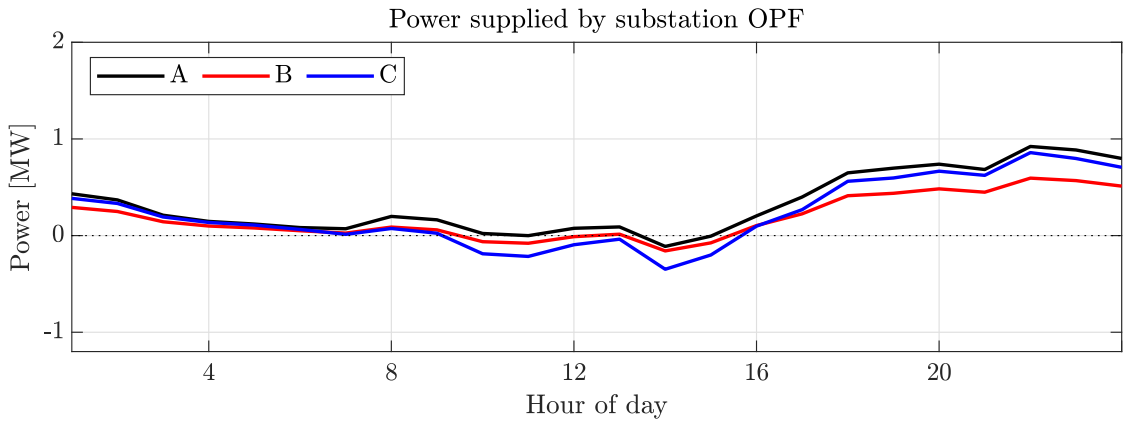



Figure 3.17: Power supplied by substation with default loads, EVs and solar PV generators, calculated by the OPF.

In the case of the active power, a breakdown of how the power is distributed is needed to fully understand the converters operation and the exchange of power between the different buses.

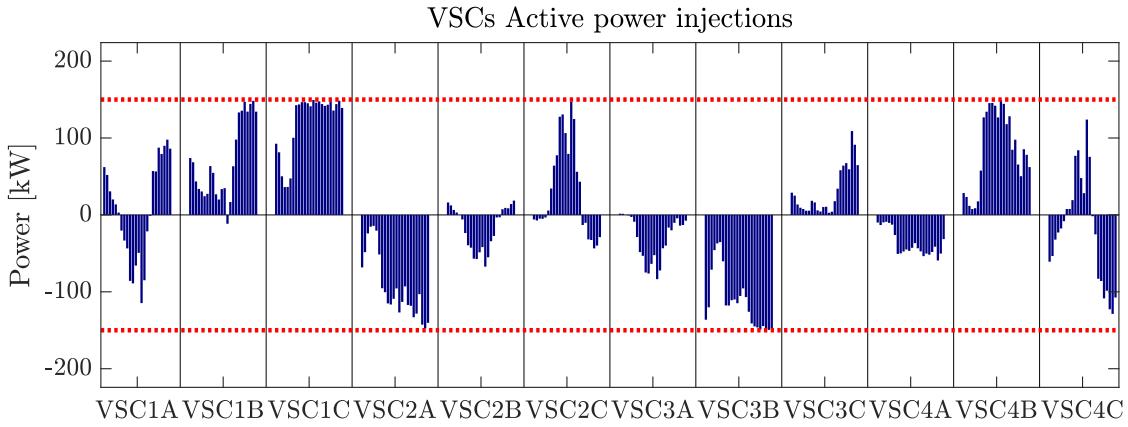



Figure 3.18: Active power injected or absorbed by converters set out sequentially for the 24hr period. The red dashed line indicates the maximum active power that the converters can inject.

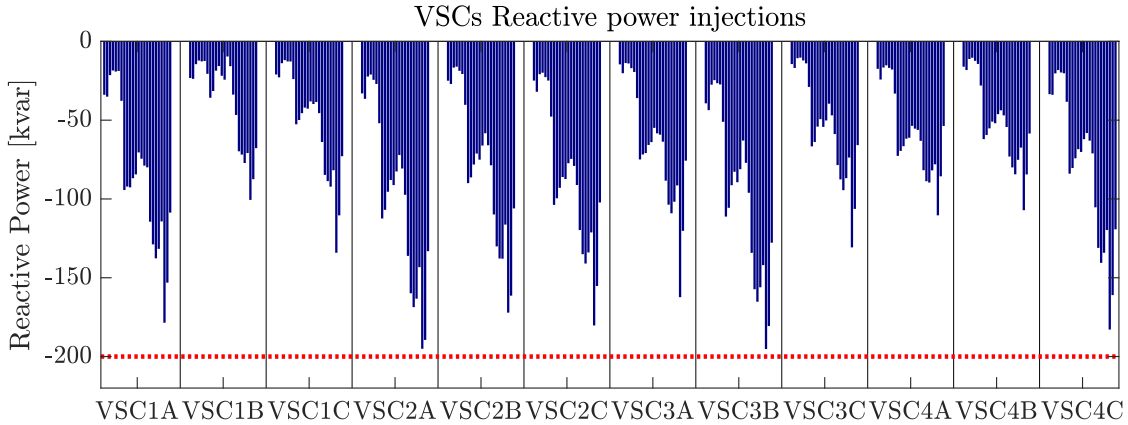



Figure 3.19: Reactive power injected or absorbed by converters set out sequentially for the 24hr period. The red dashed line indicates the maximum reactive power that the converters can inject.

Fig. 3.20 shows the internal exchange of power between the phases of each of the converters. A positive value indicates that power is being sent from that phase to the rest, whereas a negative value means that

power is being injected into that phase from the others. For instance, in the case of VSC 3, between 8 and 16 hrs power is being sent from phase C into phase A and B. However starting at approximately 17:30 hrs, the trend is changed and power is instead sent from phase B into phase C and A. A similar pattern is observed in the rest of the converters. Overall, power is transferred from phase C into phase A and B to balance out the excess solar PV generation in phase C.

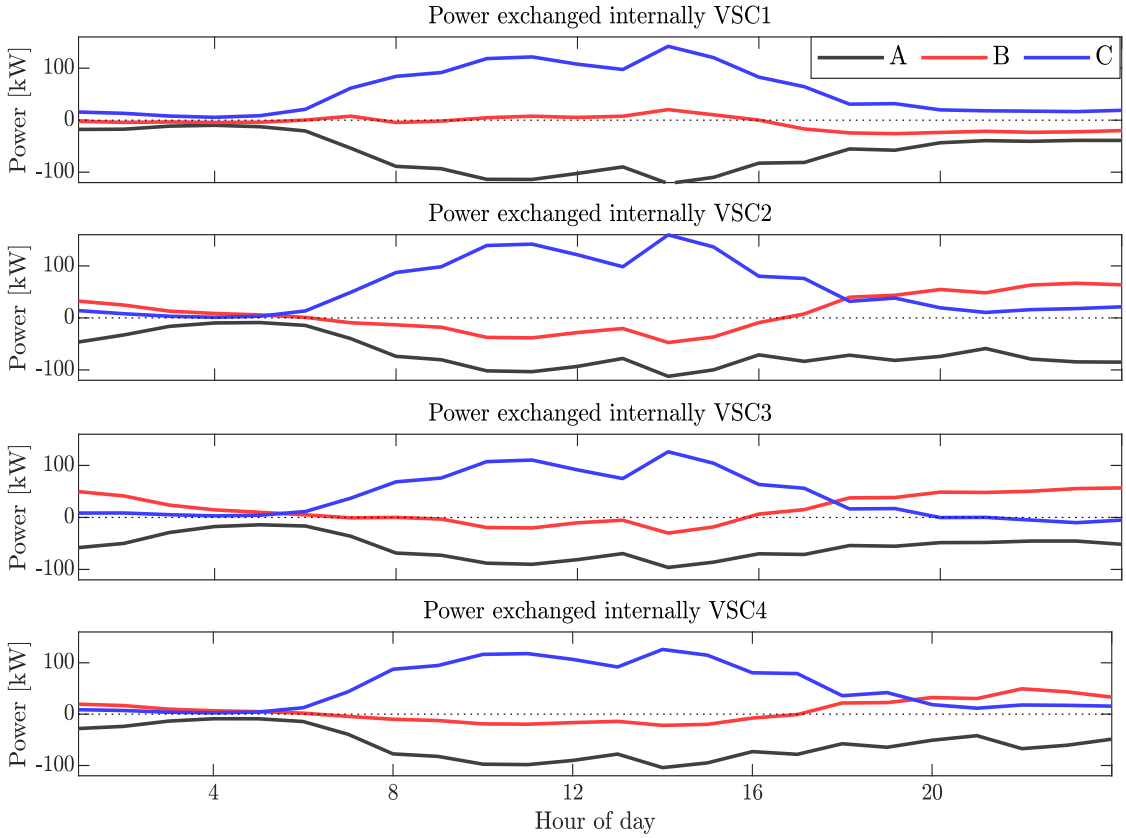



Figure 3.20: Active power exchanged internally, between the phases of each converter.

On the other hand, the exchange of power between the converters at different buses is shown in Fig. 3.21. It can be seen that for most of the day power is transferred from VSC 1 at bus 39 to VSC 2 at bus 66, from VSC 3 to VSC 4, at buses 151 and 300 respectively, and from VSC 2 to VSC 3. At approximately 14 hrs. the trend is reversed for a short period of time before increasing to a maximum during the evening hours. This could be due to the generation from solar PV and then due to the additional load from the EVs. The enhanced distribution of the power between the lines, the improved phase balance and the reactive power injections lead to reduced power losses compared to the base case, as seen by comparing Fig. 3.11 and Fig. 3.22. The total losses of the lines and transformer for the day are 1355.80 kWh, a 30% reduction from the base case.

To enable a high-level control for an ANM scheme in real-time, the OPF is needs to calculate the power set points as quickly as possible. Thus, one of the main characteristics to be evaluated is its computational tractability, especially as the AC/DC OPF incorporates more constraints due to the converters and the DC connections.

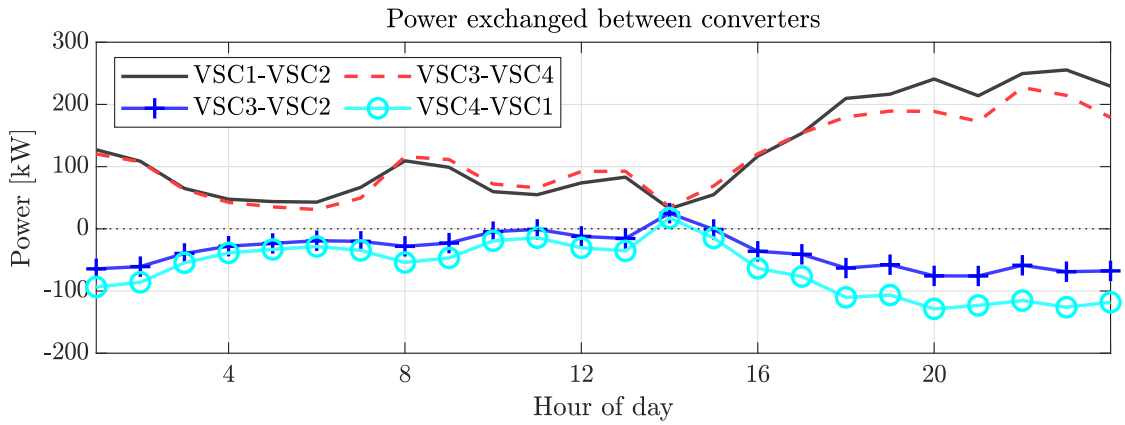



Figure 3.21: Active power exchanged between converters.

The default algorithm options in MATLAB only allow the *fmincon* IPM method to run sequentially, using only one of the cores of the computer used. To improve the scalability and speed of the algorithm, the parallel toolbox from MATLAB is used to split the algorithm and evaluate the OPF constraints in parallel, using all the available cores. This way each core can evaluate a different setpoint.

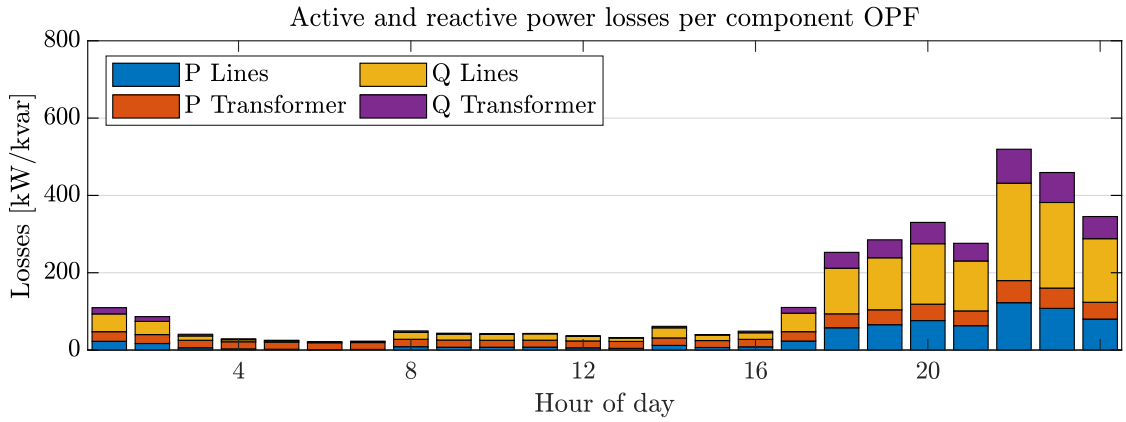



Figure 3.22: Active and reactive power losses by component with default loads, EVs and solar PV generators, as calculated by OPF.

For this particular test feeder, solving an OPF for each individual time step requires on average 365 seconds, a little more than 6 minutes, with the optimisation in parallel mode. Thus, optimising the whole day of operation in hourly steps requires a total of 146 minutes, or close to 2.5 hours, a significant increase compared to the Power Flow runtime shown in Fig. 3.12. The average serial runtime when using only one core was almost twice as much, thus the parallel option was used for the case study and only those results are reported. The OPF runtime for each time step and the iterations required are illustrated in Fig. 3.23. It can be seen that some particular time steps take longer than the average. This could be caused by the constraints, which make it harder to converge for some time steps and also due to the fact that some time steps require a larger power injection, which results in a larger number of iterations needed to reach the optimal power injection.

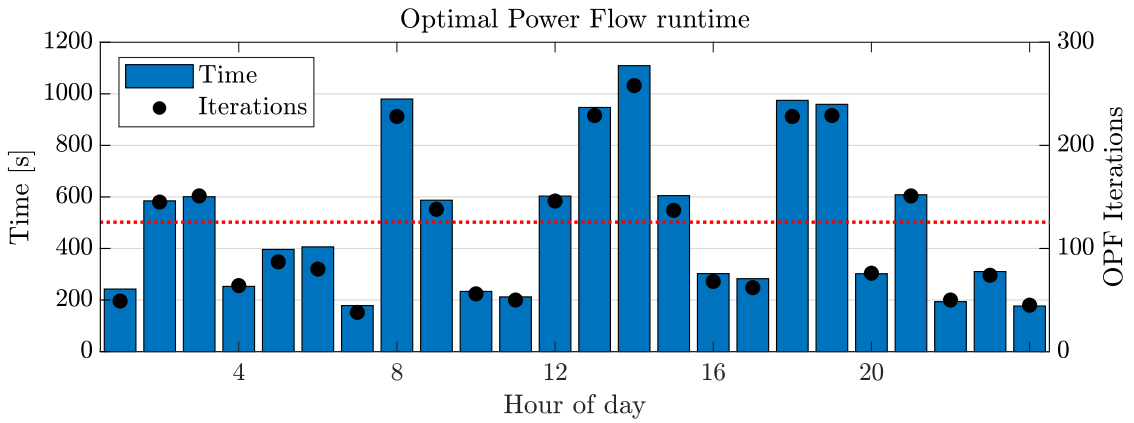



Figure 3.23: Optimal power flow runtime for each time step, the dashed line represents the average runtime. The black solid circles represents the number of iterations required by each time step to converge.

The evolution of the objective function for each time step is shown in Fig. 3.24, where the power losses have been normalised to make the comparison easier. It can be observed that all the time steps follow a similar trend, starting at the base power losses then slowly decreasing with the iterations. It is worth noting that each time step requires a different number of iterations to converge and some of them increase before decreasing again. This is because the non-optimal power injections by the converters can actually worsen the power losses.

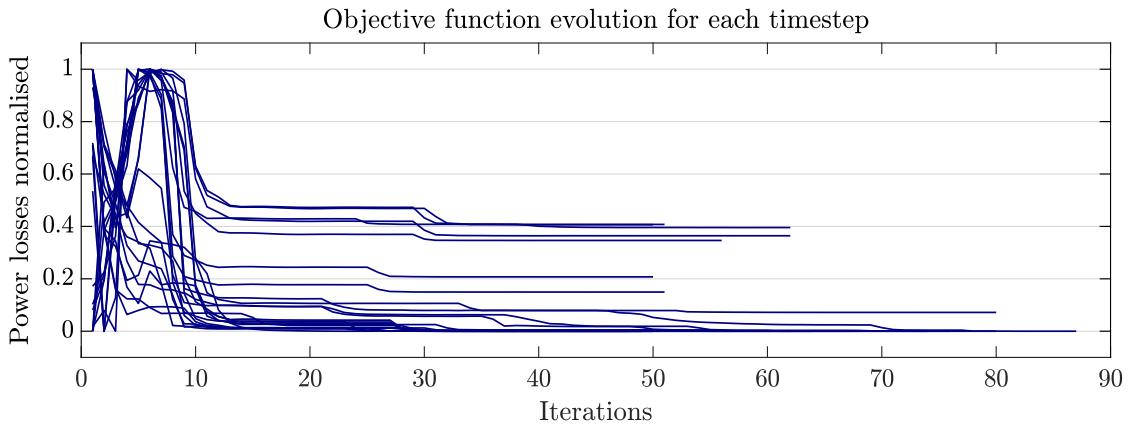



Figure 3.24: Optimal power flow evolution of objective function for each time step.

Fig. 3.25 makes a comparison of the value of the objective function (minimisation of the power losses) at the first and at the last iterations of the OPF, for each time step. From this it can be seen that although the difference for some time steps are small, especially those with low loading conditions, overall the power losses are successfully reduced.

To understand further how the optimisation progresses, a sample of how the converters power injections evolve with the OPF iterations is shown in Fig. 3.26 and Fig. 3.27 for two time steps, hour 7 and hour 14 respectively. It can be observed that some converters converge to a larger power injection than the initial, whereas others shift the direction of the initial power injection.

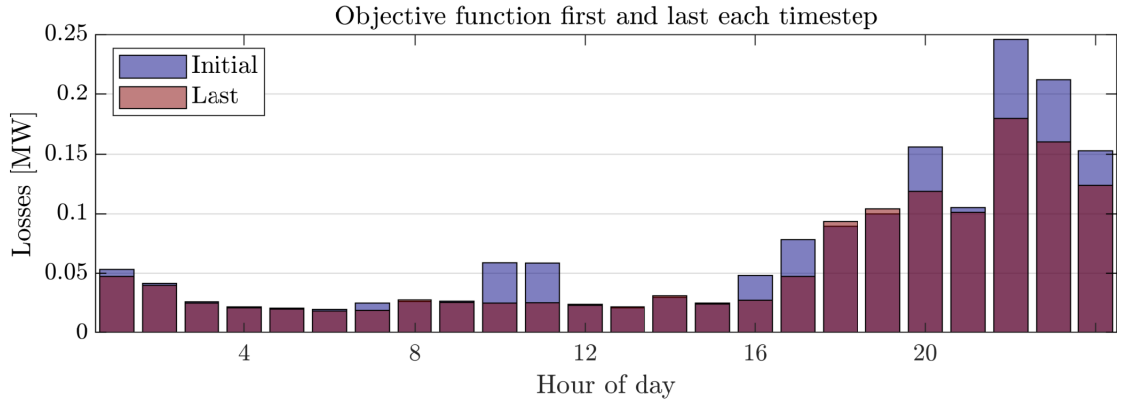



Figure 3.25: Optimal power flow evolution of objective function for each time step.

As seen in Fig. 3.23, hour 14 required the longest time to converge (1100 seconds), and it took the most iterations, whereas hour 7 required the shortest time (178 seconds), and the least iterations.

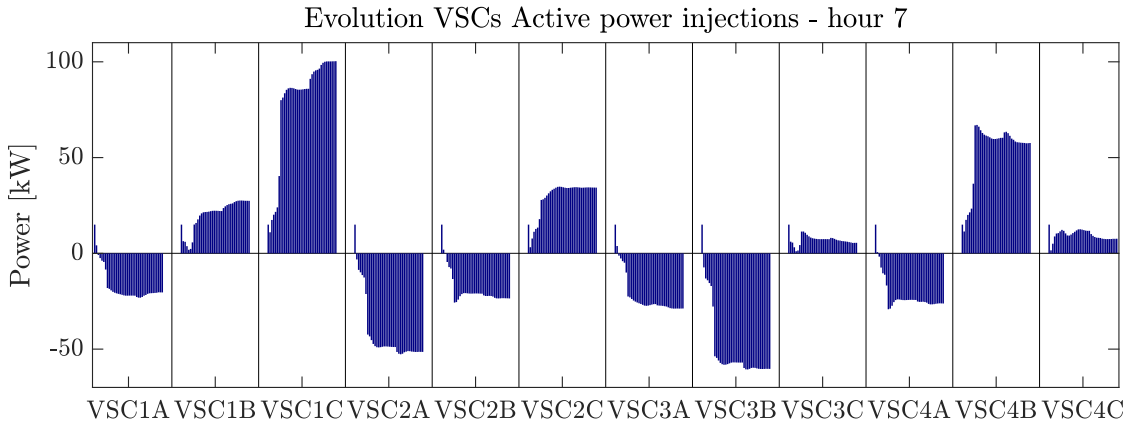



Figure 3.26: Optimal power flow evolution of the converters' power injections for hour 7. The power injections are set sequentially following the number of iterations required for the OPF to converge.

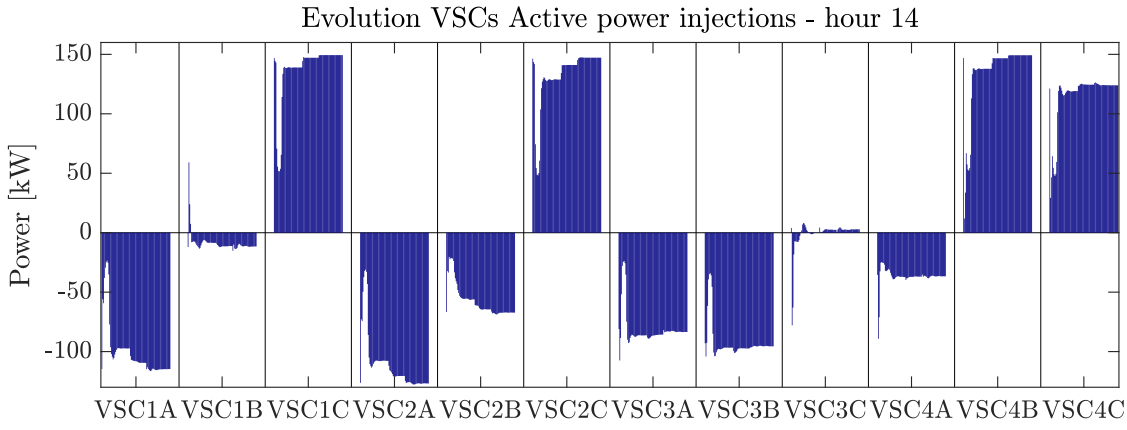



Figure 3.27: Optimal power flow evolution of the converters' power injections for hour 14. The power injections are set sequentially following the number of iterations required.

From the figures Fig. 3.27 and Fig. 3.26 it can be seen that one of the reasons for a slower convergence would be the larger power injection required for the OPF to converge. Whilst for hour 5, the maximum

injection is close to 40 kW, for hour 17 the largest injection is closer to 150 kW.

### 3.4.2 Lotherton and Canefields LV Substations

As it was discussed in section 3.1, ANM can be achieved by meshing a network by opening/closing its NOP [146, 147] or by using controllable circuit breakers as proposed in [55]. However the advantages and limitations of this ANM scheme versus the ANM using power converters are not clear, both in terms of operational benefits, long-term viability and also regarding control complexity. Hence, this case study makes a comparison between an ANM scheme via switching and ANM using power electronics converters in terms of power losses and overall network efficiency. It also demonstrates the OPF flexibility to adapt its objective function to a particular constraint, such as the loading equalisation between two transformers. It uses high-resolution data from two adjacent substations part of the OpenLV project where the ANM scheme through switching was trialled [55, 119]. During the OpenLV project, the two substations, Lotherton and Canefields, were meshed together via their NOP and using ALVIN Reclose™ devices as illustrated in Fig. 3.28, repeated from section 3.1 for clarity.

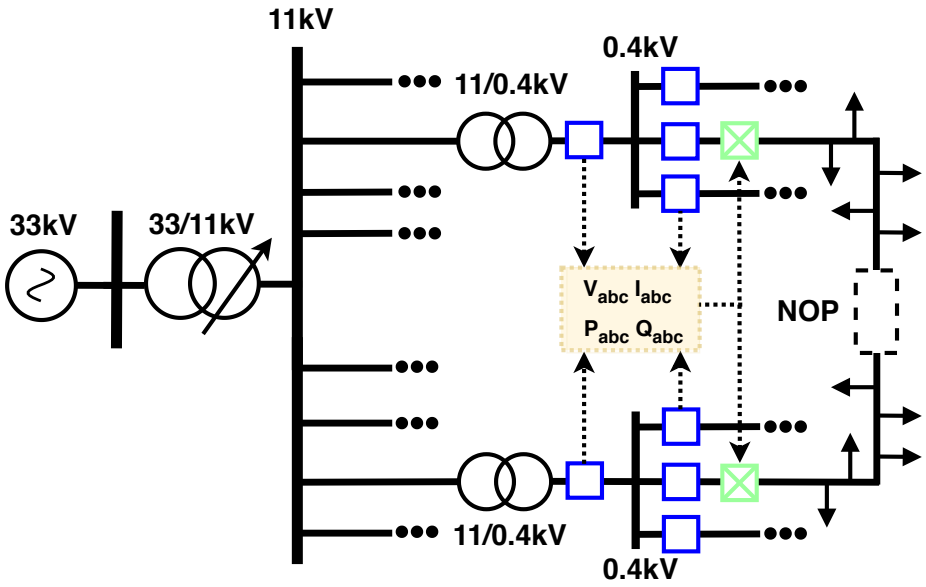



Figure 3.28: Overview of network meshing in the OpenLV project. The monitoring units are shown in blue and the ALVIN Reclose™ devices used for ANM are in green [55] (repeated here for clarity).

It is worth pointing out that the selection of the pair of substations to be connected by the ALVIN Reclose™ units was primarily based on whether they were part of the same primary substation upstream and on the permitted fault current levels [55]. Some of these limitations could be relaxed when using power converters, as back-to-back AC/DC converters could connect networks regardless of their phase-angle and voltage amplitude differences [62]. Also, as it was discussed in 1.4, the converters do not contribute to the

fault levels of the network. Nonetheless, suitable control and protection strategies would still be needed in case of faults to ensure that connections between primary circuits do not result in back-feeding.

The ALVIN Reclose™ devices at Canefields substation operate based on the loading measurements of the transformers and a forecasting algorithm [55]. Whereas the ALVIN Reclose™ devices at Lotherton remain closed. The ANM scheme is enabled by peer-to-peer communication between the monitoring devices at both substations [55]. A network diagram of the actual substations and underground cabling was obtained from the Western Power Distribution data portal, and is shown in Fig. 3.29 [148]. This image shows the LV underground cables that comprise the feeders and also the service cables used to supply the customers. The corresponding satellite image of the location of both substations and the NOP can be seen in Fig. 3.30, where the properties supplied are observed more clearly. It is evident that the majority of the loads are residential properties and some of them have solar PV installed. Additionally, both maps are overlaid in Fig. 3.31 to make the visualisation easier [148].

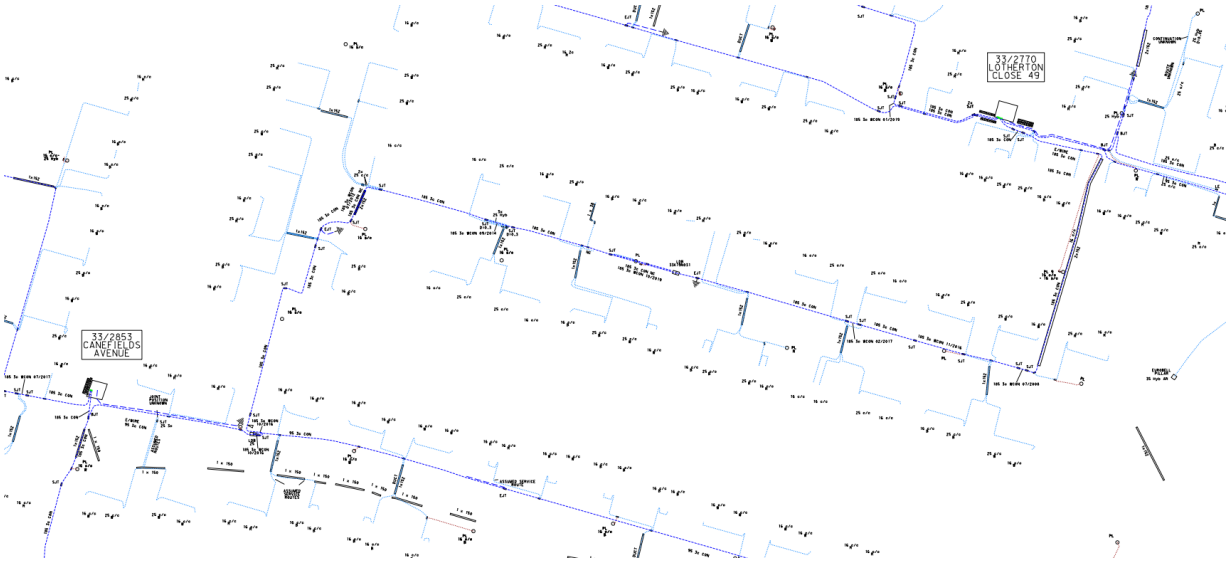



Figure 3.29: Network map image of Lotherton and Canefields substations. The square on the upper right corner and on the bottom left corner show where the LV substations are located. The blue lines depict the LV and service cables used to supply the customer properties. Reproduced with permission from Western Power Distribution [148].

The active and reactive power measurements can be used to calculate the total load supplied by each substation transformer, the available capacity and the level of phase imbalance in the feeders. Historic data with a 1-minute resolution is available for the active and reactive power, the voltage and the current for each phase of the outgoing feeders. The same is available for the main substation busbar. It is worth noting that this data represents the aggregated load of all the customers connected to the substations. Individual load data is not available to the network operator and thus not used in this case study. A comparison of the total power supplied by both substations before the network meshing was in operation is depicted in Fig. 3.32, where it can be seen that Lotherton has a much higher loading throughout the day compared to Canefields.




Figure 3.30: Aerial image of Lotherton and Canefields substations. The yellow line illustrates the path followed by the underground feeders of each substation and the red square illustrates the location of the NOP.

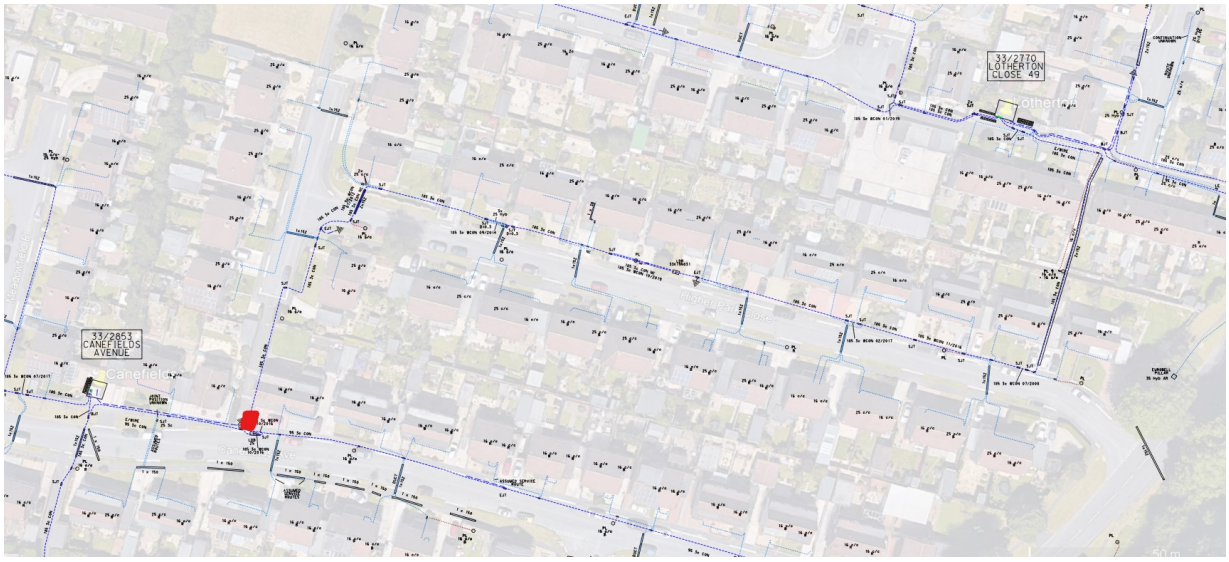



Figure 3.31: Aerial image of Lotherton and Canefields substations with the network map overlaid. The blue lines depict the LV and service cables used to supply the customer properties and the red square illustrates the location of the NOP. Reproduced with permission from Western Power Distribution [148].

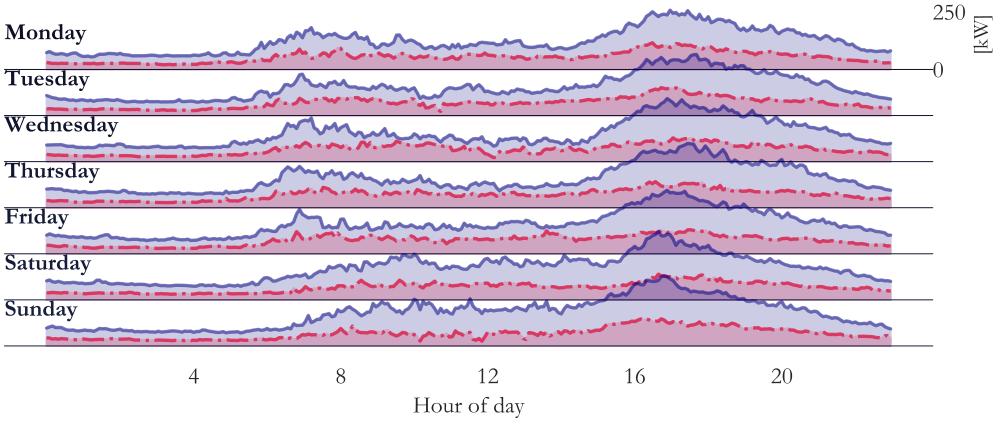



Figure 3.32: Total active power supplied by each substation, Lotherton (blue) and Canefields (red), over a sample week in November.

Also, the phase imbalance is significant in both substations, as seen in Fig. 3.33 for Lotherton and Fig. 3.34 for Canefields. The neutral current is particularly large in Lotherton substation, with an average peak of almost 300 A, due to the difference between the loading of phases A and C compared to B. Canefields substation displays a lower phase imbalance, however during the hours of peak load the neutral current still represents around one third of the most loaded phase.

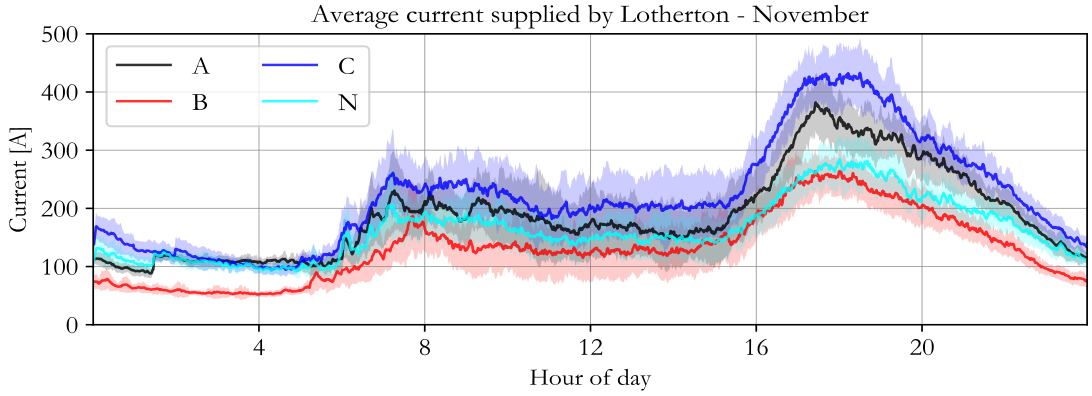



Figure 3.33: Current supplied by Lotherton substation for November. The solid line depicts the average for that month and the shaded area represents the standard deviation.

The status and power going through the ALVIN Reclose™ units is also recorded in the data. The total power transferred and the hours at which the ALVIN Reclose™ units opened/closed are illustrated in Fig. 3.35 and Fig. 3.36 for March, once the ANM trial was operational. As it is expected, when the ALVIN Reclose™ units are closed the total active power transferred from one substation to another varies from day to day, as it depends on the impedance and the voltage of the feeders. On average, in Fig. 3.35 it can be seen that the peak active power transferred from Canefields to Lotherton is around 50 kW and occurs in the evening hours. However, there is a lot of variations between days, with some of them having a peak of approximately 80 kW.

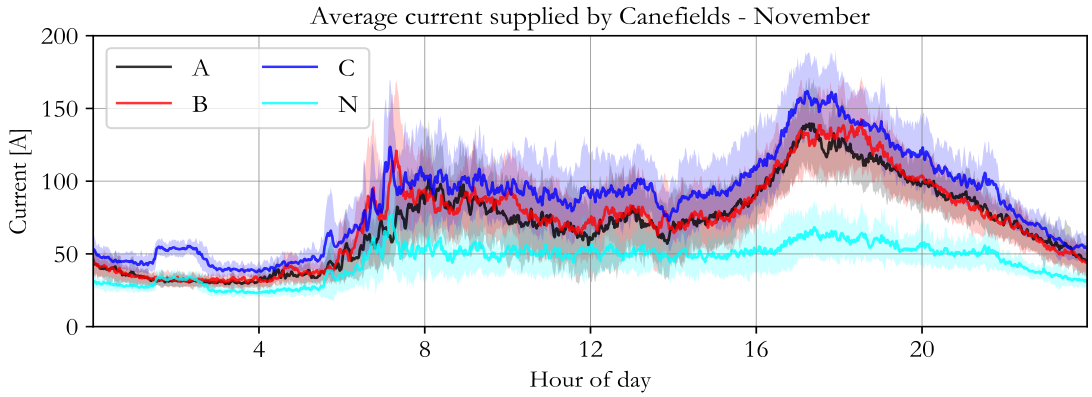



Figure 3.34: Current supplied by Canefields substation for November. The solid line depicts the average for that month and the shaded area represents the standard deviation.

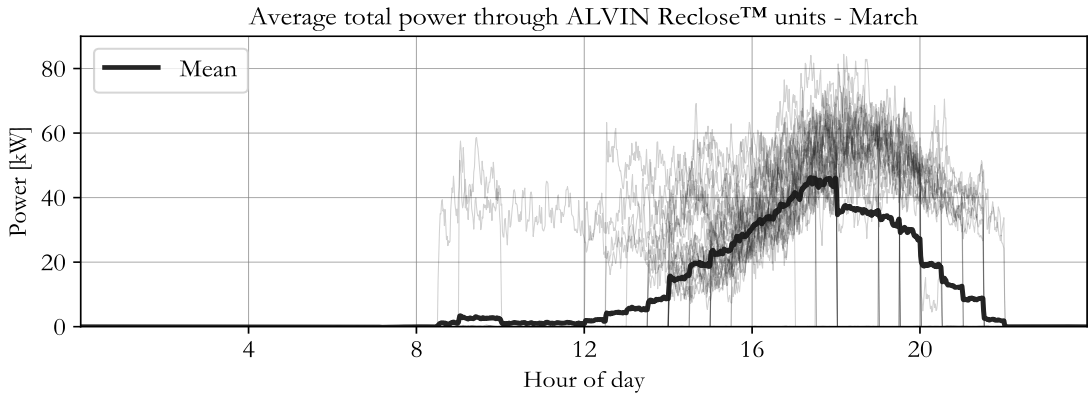



Figure 3.35: Daily samples and monthly average of total power flowing through the ALVIN Reclose™ units.

It can be seen in Fig. 3.36 that on average the ALVIN Reclose™ units remain open for the majority of the day, and then close for the afternoon-evening peak (12-22 hours). Also, it is worth noting that for some days they remain close for only a couple of hours whereas for other days the duration is longer. The historical

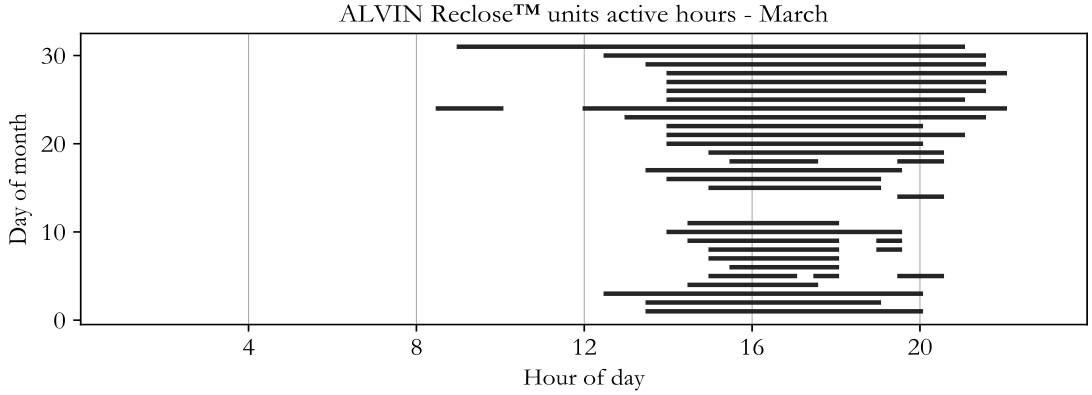



Figure 3.36: Actuation times of the ALVIN Reclose™ units are indicated by the solid line.

loading data, together with the transformer and cable impedance data, provided by the network operator, are used to build a realistic model of the substations and the head of feeders. Table 3.1 presents a summary of this data. A three-phase model of the substations is built in OpenDSS and used to obtain the baseline conditions and the power flows for the ANM via switching.

Table 3.1: Network data and peak load information.

| Substation | Transformer rating [kva] | Feeder head 1 [kW] | Feeder head 2 [kW] | Feeder head 3 [kW] | Feeder head 4 [kW] |
|---|---|---|---|---|---|
| Lotherton | 500 | 78.26 | 44.83 | 128.4 | 95.80 |
| Canefields | 315 | 124.56 | 0.0 | - | - |

The benefits of the meshing approach are then compared with the alternative based on power converters, where two converters are used to mesh the two substations, replacing the NOP, as in Fig. 3.37.

Although both ANM schemes should attempt to release network capacity by improving the network parameters, the power converters are capable of very fast adjustments, providing a higher degree of control over the power flows, as it was shown in the previous case study. According to the average power transferred through the ALVIN Reclose™ units, as shown in Fig. 3.35, two three-phase converters, each with a rated power of 90 kVA, will be used to connect the two substations via their NOP as shown in Fig. 3.37.

The AC/DC OPF is used to calculate the converters optimal power injections, whereas OpenDSS is used to calculate the power flow with the ALVIN Reclose™ units. In this case, the objective function of the OPF is to achieve loading equalisation between the two transformers, rather than minimising the power losses, in order to make a direct comparison to the switching scheme implemented in the OpenLV project.

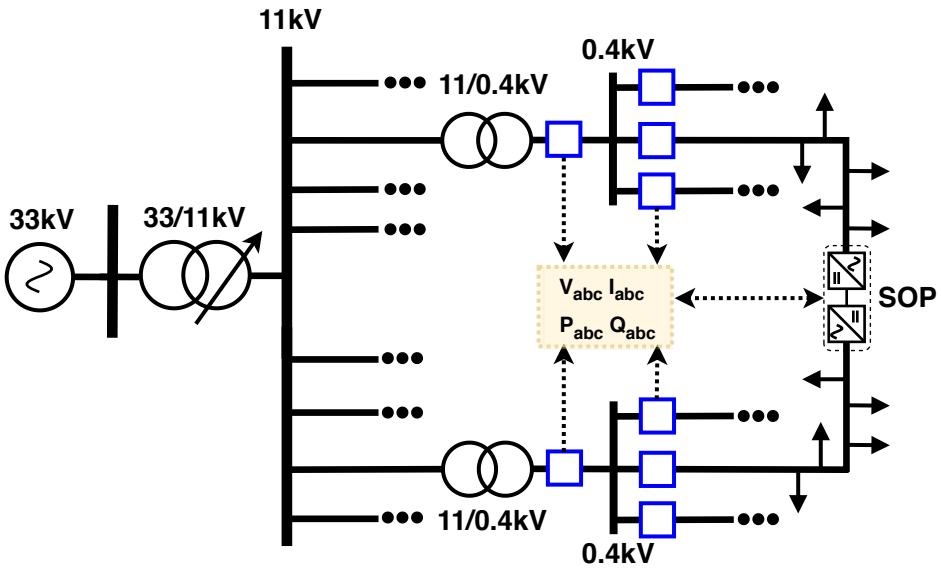



Figure 3.37: Two LV substations connected using power converters, with the monitoring units shown in blue. This back-to-back connection between the converters is also referred to as SOP (figure repeated here for clarity).

The aggregate profile of the load supplied per-phase of each individual feeder is obtained by averaging the loading data for the same weekday over a month period. This is used to model the power supplied by the substation. Also, the average opening/closing hours are used for the ALVIN Reclose™ units.

As can be seen in Fig. 3.38, both network management schemes keep the loading of the transformers almost equalised compared to the base case. Load is transferred from the Lotherton substation to the Canefields substation and overall a lower utilisation of the transformers is achieved. However, the power converters can operate for the whole day whereas the switching occurs only at specific times, the average of those shown in Fig. 3.36.

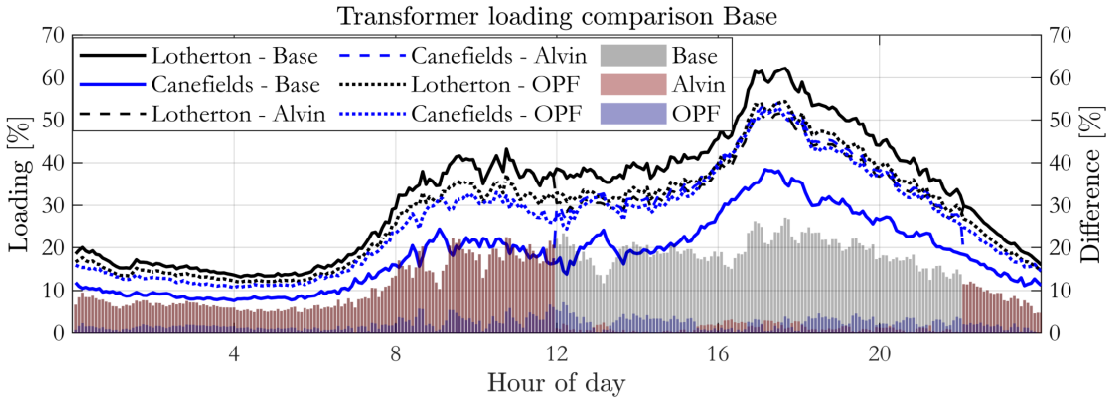



Figure 3.38: Comparison of the transformers loading and their difference. The solid lines indicate the results of the base case, whereas the dashed and dotted lines indicate the results of the switching and OPF, respectively. The bars indicate the loading difference between the transformers.

Moreover, a large difference is observed in the current transferred between the substations via the ALVIN

Reclose™ units and the converters, both in terms of magnitude and phase imbalance. The converters regulate the power going into each of the phases, which helps reduce the phase imbalance, whereas the ALVIN Reclose™ units simply provide an alternate path for the power flows. This results in a much lower neutral current for both feeders, as can be seen in Fig. 3.39. In the case with the converters, the neutral current is reduced 40% from the base case. Also, the magnitude of the current transferred is reduced, which translates into reduced power losses.

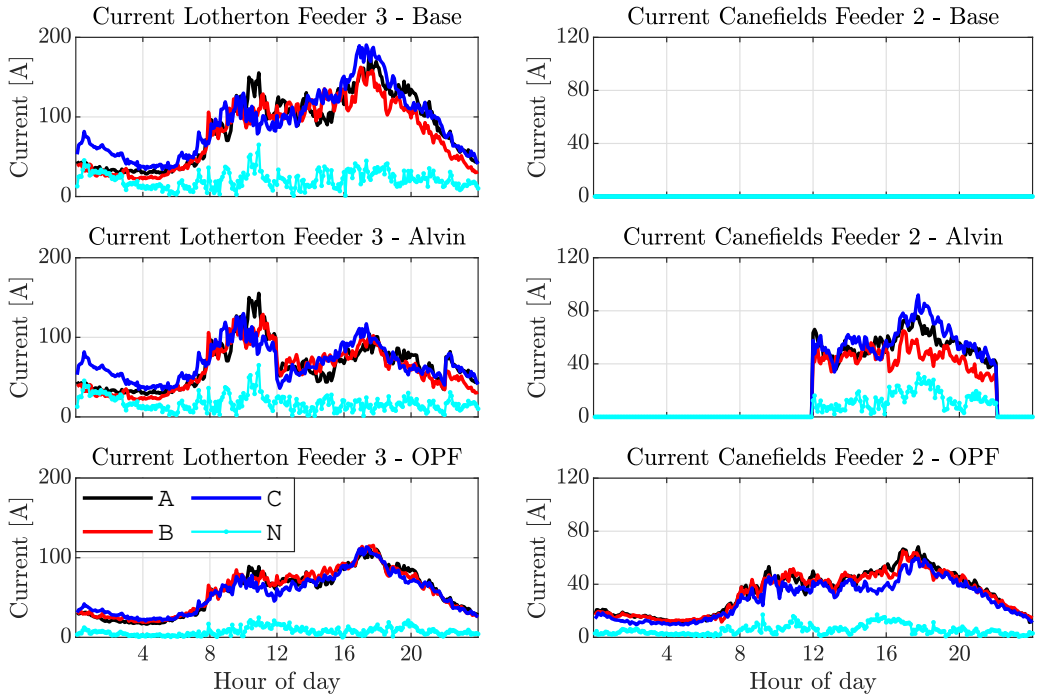



Figure 3.39: Comparison of the current flowing through Feeder 3 and Feeder 2 for the base case, with the ALVIN Reclose™ units and with the converters.

To assess the long-term viability of each network management approach, both schemes are compared when different volumes of LCTs are connected to the substations. In particular, the schemes are compared when the LCTs are connected at only one substation or at both. The scenarios analysed are described in Table 3.2. Fig. 3.40 shows that the goal of equalising the load between the transformers (Lotherton Tx1 and Canefields Tx2) is achieved with both schemes for all the cases. Compared to the base case, the maximum loading of transformer 1 is reduced significantly with both ANM schmes, whereas the maximum loading at transformer 2 (Tx2) is increased to a comparable level to transformer 1 (Tx1). Nonetheless, the converters keep a lower overall loading, as seen in cases 5-7, where the total kVA capacity released at peak hours is approximately 10-30 kVA higher for the OPF case compared to the ANM through switching.

This is in part because for some cases the optimal distribution of the power might not be achieved by switching, where the power flow is determined by the network impedance and cannot be controlled. This is illustrated further When looking at the different cases, where it can be seen that the power losses are affected by the specific location of the LCTs.

Table 3.2: Cases with increased LCTs penetration.

| Case | Lotherton EVs [kW] | Canefields EVs [kW] | Lotherton solar PV [kW] | Canefields solar PV [kW] |
|---|---|---|---|---|
| 1 | - | 66 | 100 | - |
| 2 | - | 99 | 130 | - |
| 3 | 66 | - | - | 100 |
| 4 | 99 | - | - | 130 |
| 5 | 33 | 33 | 50 | 50 |
| 6 | 47.3 | 51.7 | 65 | 65 |
| 7 | 66 | 66 | 100 | 100 |

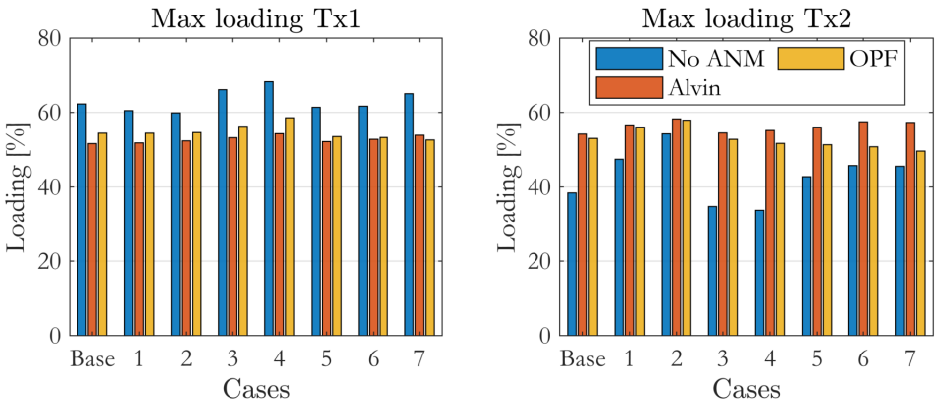



Figure 3.40: Comparison of transformers loading for each case.

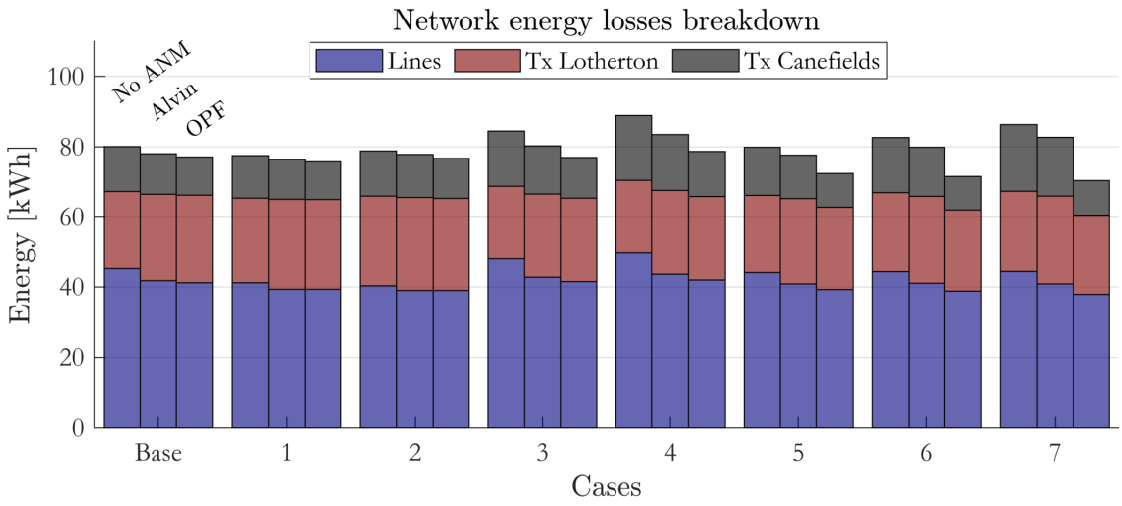



Figure 3.41: Comparison of daily energy losses by network component. From left to right, the bars indicate the case without ANM first, followed by the Alvin case and the OPF.

For example, case 2 and case 4 both have the same penetration level of LCTs. However the losses are higher when the EVs are connected to Lotherton and the solar PV to Canefields. Overall, the energy losses are 15-20% lower for the ANM using the converters compared to the case with the ALVIN Reclose™ units.

It is worth highlighting that a better distribution of the power from the PV generation can help minimise the amount of reverse power flowing to the substation and also to reduce the overall distance that the power needs to travel.

These results indicate that despite the much simpler control needed for the ANM with ALVIN Reclose™ units, its performance is limited by the open/close switch operations with no control of the power flow. The ANM with power converters has a more flexible control scheme that can be adapted as the level of LCTs increases. This is particularly important because the DNOs have little control over where customers install LCTs. Having the option of continuously controlling the power flows translates into lower currents with a lower phase imbalance, it also enables specific control objectives, such as loading equalisation or losses minimisation.

### 3.4.3 European LV Network

To examine more in detail the advantages of VSCs and MTDC connections, a more realistic LV network will be used in the following case. A portion of the IEEE European Low Voltage (EULV) network from [143,144] is used, which is representative of LV networks in the UK as described in section 2.1. In this case, the topology and connectivity of the network is available as can be seen in [143, 144]. This network consists of one feeder connected to a main substation with a 315 kVA transformer. The rated voltage level is 416 V (phase-to-phase) and there are a total of 135 three-phase lines (underground cables) with 65 loads (each with a peak power between 3-5.5 kW) distributed randomly across the network, as illustrated in Fig. 3.42 [143, 144]. The cable parameters are given as positive and zero sequence impedances and length. Detailed information about the modified European network is presented in Table 3.3. Additionally, 35 EVs, each with a demand of 3.3 kW, were assigned to some of the loads. Similarly, 20 solar PV installations between 1.5-3.5 kW were randomly assigned. Both LCTs were assigned a power factor of 1.

Table 3.3: Details of the modified European LV network.

| Number of buses | Feeder length [m] | Number of lines | Number of loads | Average power factor | Phase A % | Phase B % | Phase C % |
|---|---|---|---|---|---|---|---|
| 136 | 464.3 | 300 | 65 | 0.95 | 43.5 | 38.5 | 18 |

Two configurations for the power converters will be analysed and compared. The possible location of the power converters is highlighted by the red squares in Fig. 3.42. The first case will examine the benefits

of having individual power converters (phase-balancers) at the end of the feeders. These phase-balancers can only inject power between the phases of the same feeder via its DC bus but cannot transfer power between feeders. This case is useful whenever the geographical location is not optimal to join two or more feeders or when the feeders do not have a complimentary load/generation profile. Also, a phase-balancer approach could be tailored to the feeder in question, sizing the converter according to the level of unbalance and the peak load of the feeder. The second case will consider the AC/DC connection between neighbouring points of the feeder, those that are close to each other as shown in the diagram in Fig. 3.42. This will provide alternative paths for the power flow, with the goal of improving the operation of the network.

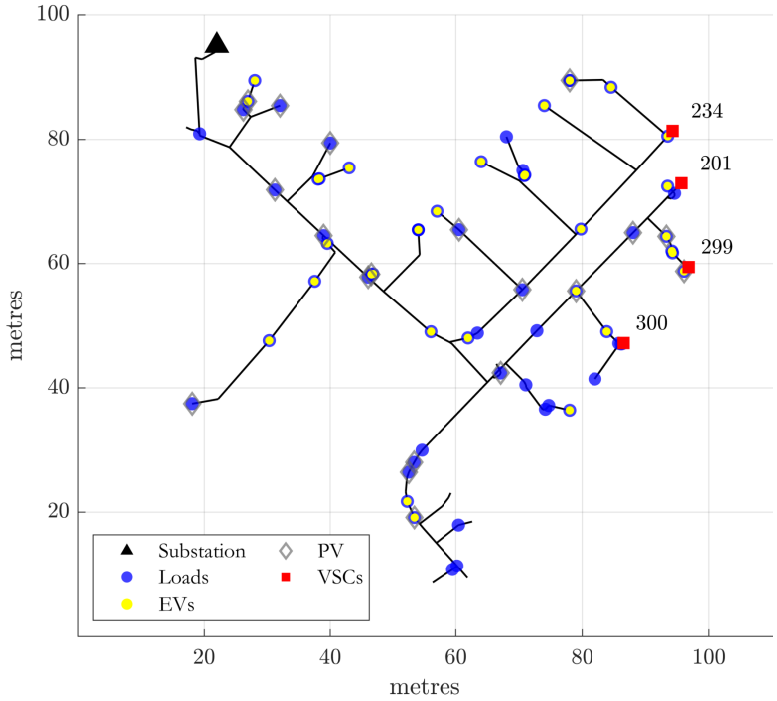



Figure 3.42: Diagram of the European low voltage network. The substation is highlighted by a black triangle, whereas the loads are depicted by the blue circles. Those customers with EVs have a yellow circle overlayed and those with solar PV have a diamond marker. The potential location of the VSCs is illustrated by the red squares [143, 144].

A representative 24 hour load and generation profile is used, similar to the IEEE 123-bus test feeder case study in 3.4.1. Due to the size of the network, it is expected that the both the power flow and the OPF will take longer to converge. The results of the base case without power converters during peak power conditions (22 h) and the total daily energy losses are summarised in Table 3.4. The results reported for energy losses are the total losses for the 24 hour period modelled.

Table 3.4: Base case results European LV network for the daily load profile.

| **Line losses [kWh]** | **Transformer losses [kWh]** | **Total energy losses [kWh]** | **Transformer loading max [%]** | **Vmin [pu]** |
|---|---|---|---|---|
| 110.0 | 49.3 | 159.3 | 116.7 | 0.892 |

Fig. 3.43 shows the average time taken by the power flow to converge and also the runtime at each time step. It can be observed that the average runtime is similar to that of the IEEE 123-bus test feeder, despite the additional number of lines and nodes.

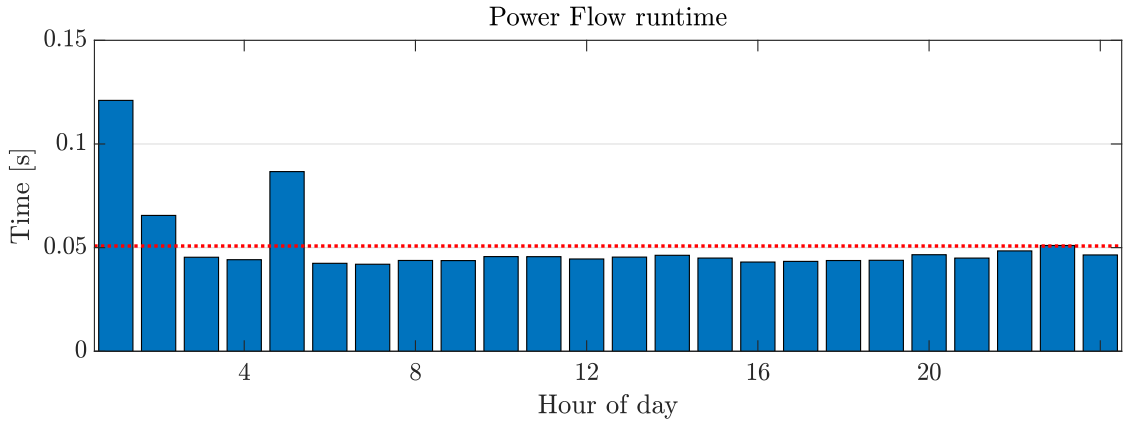



Figure 3.43: Run-time of the power flow for each time-step, the dashed line represents the average runtime.

The current supplied by the substation is shown in Fig. 3.44, where a large phase imbalance is observed. In particular, it can be seen how the neutral current has a higher value than phase C during the hours of peak load.

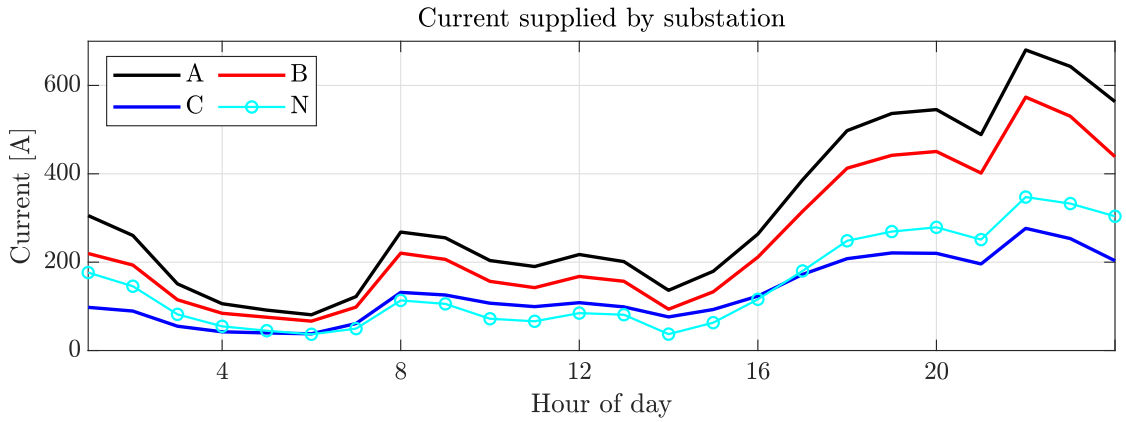



Figure 3.44: Current supplied by substation with default loads, EVs and solar PV generators.

Table 3.5 compares the results for the base case, the case using phase balancers and the case with DC connections between the converters. It can be seen that in this case, the difference between having DC connections between the converters is minimum. However, the power losses are reduced in each case by decreasing

the phase imbalance and the overall peak current, as shown in Fig. 3.46 for the case with DC connections between the converters.

Table 3.5: Results comparison European LV network.

| Case | Line losses [kWh] | Transformer losses [kWh] | Total energy losses [kWh] | Transformer loading max [%] | Vmin [pu] |
|---|---|---|---|---|---|
| Base | 110.0 | 49.3 | 159.3 | 116.7 | 0.892 |
| Balancers | 96.7 | 46.8 | 173.94 | 114.0 | 0.929 |
| MTDC | 96.1 | 46.9 | 184.1 | 114.5 | 0.929 |

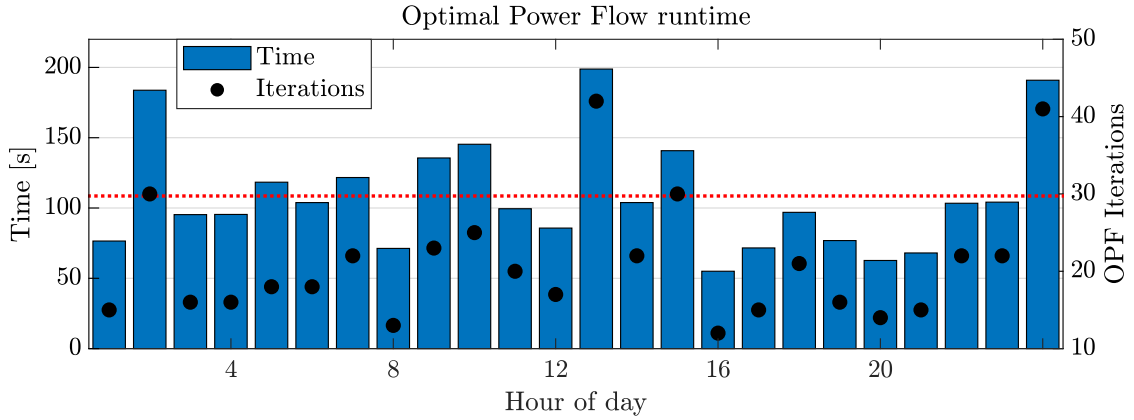



Figure 3.45: Optimal power flow runtime for each time step, the dashed line represents the average runtime. The black solid circles represents the number of iterations required by each time step to converge.

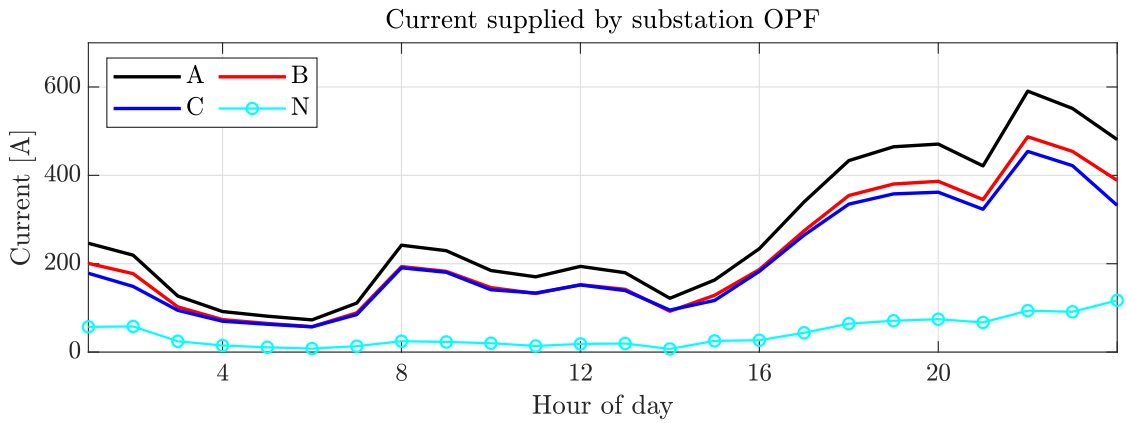



Figure 3.46: Current supplied by substation with default loads, EVs and solar PV generators, calculated by the OPF.

Also, the respective evolution of the objective function for each time step is shown in Fig. 3.47. Similar to the IEEE 123 case, it can be observed that all the time steps follow a decreasing trend after approximately

10 iterations, ultimately settling in the optimal point of minimum power losses. The active power injections of the converters are displayed in Fig. 3.48.

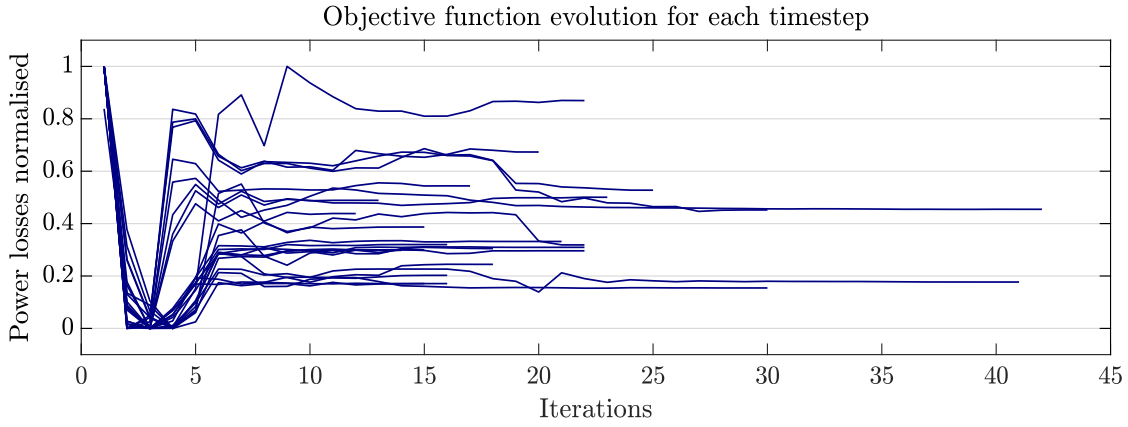



Figure 3.47: Optimal power flow evolution of objective function for each time step.

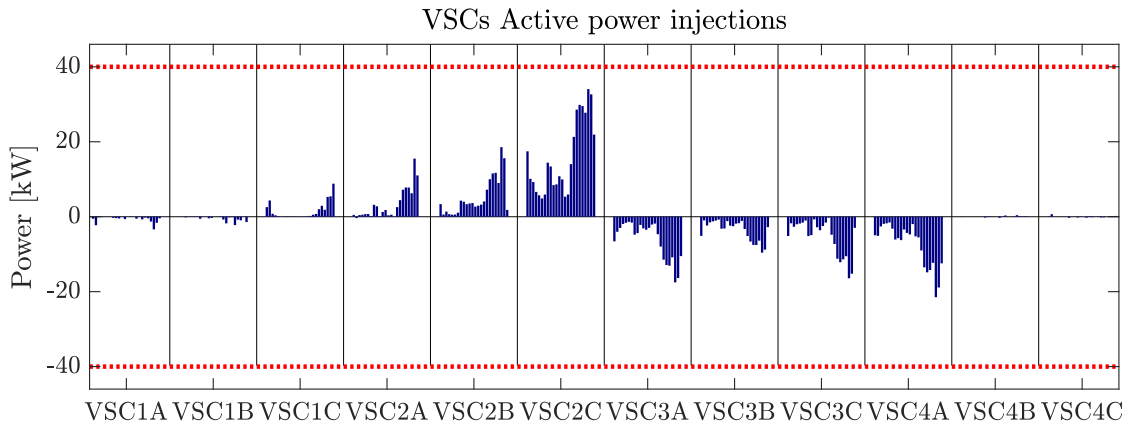



Figure 3.48: Active power injected or absorbed by converters set out sequentially for the 24hr period. The red dashed line indicates the maximum active power that the converters can inject.

## 3.5 Limitations of the OPF for ANM control

From the results presented in the case studies, it is clear that an ANM scheme that has power electronic devices at its disposal can brings important benefits over ones without, by making available more of the latent network capacity. The high-level control proposed for those studies was based on an AC/DC OPF and was an extension of the original OPF formulation from 2.3. However, there are some characteristics that prevent the OPF from being suitable for real-time operation as part of an ANM scheme and this needs further attention. In spite of the good convergence properties and the adequate treatment of the non-linear characteristics of the OPF, the IPM, and other widely used deterministic algorithms used as solvers, have a significant drawback: they cannot guarantee that the solution reached is a global optimum for this class of problems [102]. That is, the power set points obtained for the power converters through the OPF might correspond to a local

minimum and there might exist a solution that yields better results. One of the main factors responsible for this is the intrinsic non-convexity of the OPF problem. Particularly, the power/current injections and the inequalities that enforce network constraints, such as the voltage limits, as presented in sections 2.2 and 2.3. This non-convexity also results in a large computational burden as the size of the distribution network increases, particularly as phase imbalance is being considered. All of this make the OPF unsuitable for realistic LV networks such as the ones presented in the case studies. Additionally, as highlighted in section 2.2, LV feeders are composed of multiple short lines with different loads at each phase of each bus. This makes it hard to aggregate (lump) or average portions of the network without significant loss of accuracy. To overcome this obstacle, several strategies have been reported in the literature such as the use of linear approximations and the application of convex relaxations [149]. Given the characteristics of distribution networks, DC approximations do not represent the behaviour of these networks accurately enough. Hence, instead of using linear approximations for the OPF equations, recent works have proposed the use of convex relaxation methods such as semi-definite programming (SDP) and second order cone programming (SOCP) [149, 150]. Convex relaxation methods are used to surround the non-convex feasible space of the OPF problem within a larger convex space and through this the global solution can then be obtained [149].These relaxations can provide necessary conditions for determining the feasibility of the problem and they can converge to the global optimum, provided that the duality gap between the OPF and the relaxation is zero [104, 149]. Although convex relaxations present a promising alternative to the usual nonlinear optimisation algorithms, such as the IPM, there are several unresolved questions regarding their ability to solve the OPF problem for a broader range of networks. Recent studies have reported that the convergence of such methods can be severely compromised by both the physical properties of the network and by the problem formulation [151,152]. In particular, when the complexity and size of the networks starts to increase, as in the case of DNs, the convergence of these methods is not guaranteed. Also, there is an inherent problem of ensuring that the solution obtained through the relaxation is applicable to the original problem. Thus, a compromise has to be achieved between using a realistic representation of DNs and reaching a local optimum, or making several assumptions, often unrealistic, about DNs and converging to a global optimum with no guarantees that the solution is applicable to the original network. The former option is preferred in this thesis due to the nature of both DNs and LCTs.

In addition to the issues of convergence, the AC/DC OPF formulation has, so far, assumed complete observability of the network states, such as voltages at every node and current flowing in all the lines. These values are then used to calculate the optimal set points. However, as was described in section 3.1, a complete set of these values would not normally be available in the DNO control room. Usually, a DNO can install monitoring devices at their LV substation (transformer and head of feeders) and obtain measurements of the aggregated load observed along the feeders, as shown in the Lotherton and Canefields case study in sec-

tion 3.4.2. However, as was described in chapter 2 and demonstrated in this chapter, traditional optimisation methods require full observability of the network and the loads to calculate the optimal set points. Fig. 3.49 illustrates more clearly the challenge of partial observability when attempting to implement an ANM scheme. Attention also needs to be paid to the communication of calculated set points to the power converters. This becomes more challenging if the converters are not located at the substation level but downstream in the middle of or at the end of the feeder.

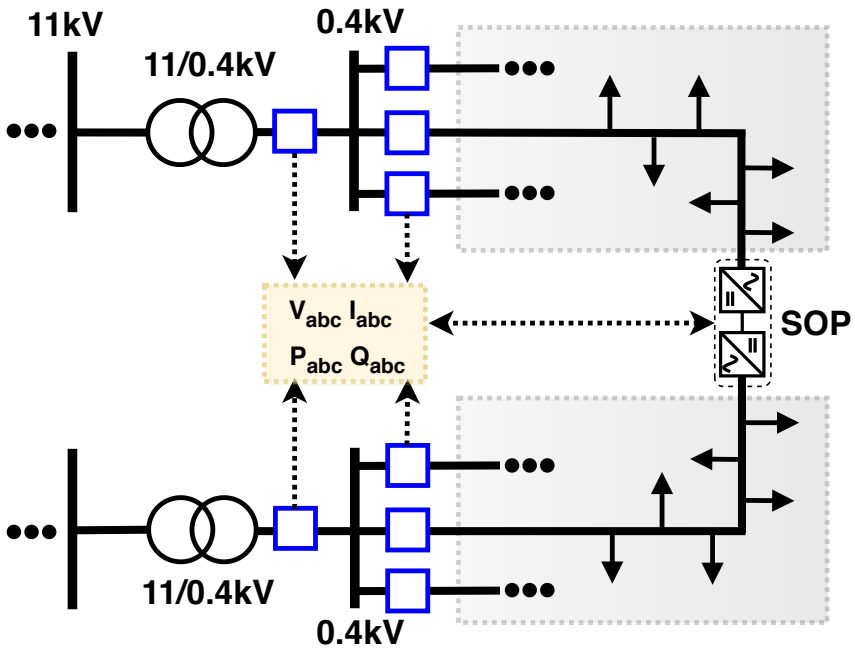



Figure 3.49: Overview of the typical monitoring available in LV networks. The grey shaded area indicates the lack of visibility of individual buses and loads along the outgoing LV feeders.

## 3.6 Alternatives to the OPF

A straightforward approach to developing the lower layers of a hierarchical control system is to create a set of simple control rules for the converters, derived from a centralised non-hierarchical algorithm. This involves a statistical analysis of the results obtained by a centralised approach, such as the OPF, over a series of scenarios. The objective is to determine the relations between the set points to be used for active and reactive power of the converters and the key parameters of the network. An example might be to look for a relationship between the voltages observed at the terminals of a particular converter and the power set points for that converter chosen by the OPF.

Another relationship that could be examined is between the set-points and the current at the head of the feeder, since these might be expected to indicate the overall phase imbalance in the network. These

relations can be identified by performing a regression analysis on the data points obtained from the OPF scenarios. Control rules could be produced from a dataset of results obtained via the AC/DC OPF. The authors in [153] proposed a two-stage control approach for FACTs devices by leveraging the accuracy of the OPF algorithm and using these results as input to a regression model. Using this approach, the authors were able to map key network measurements to the optimal set-points of the FACTs devices and produce a local controller that uses only partial network information. Both a sensitivity analysis and a brute force approach were used to select the measurements that would provide the most relevant information. To find the solution to the regression problem, the authors used the Orthogonal Matching Pursuit (OMP) algorithm. A similar study [154], used quadratic regression functions to create control rules for the reactive power produced by wind farms.

A methodology similar to the ones described here could be used to create control rules for a decentralised hierarchical ANM from a dataset of results obtained via the AC/DC OPF. The following three-phase measurements could form the basis of the dataset: the voltage at the AC terminals of the VSCs, the active and reactive power being supplied by the transformer, the value of the current flowing to the feeders and its direction (in the case of reverse power flows). The rules derived from the training dataset could then be validated using further scenarios in the same network or different networks. The identified relationships could be programmed into actual converters and the results could be compared to the centralised OPF. Finally, the AC/DC OPF and the regression could be integrated into a hierarchical control in which optimal solutions would calculated off-line for a set of scenarios, the regression coefficients would be found and functions implemented on-line to get produce the set-points for the converters in the network.

A further alternative is the conversion from a centralised OPF to a distributed optimisation algorithm, one in which the solution would be based on knowledge and communication between adjacent nodes only. Recent works have proposed several alternatives to the centralised OPF [155–157]. One of the methods that seems particularly effective for this problem is the alternating direction method of multipliers (ADMM). However, in order to solve the OPF in a distributed fashion, several assumptions need to be adopted in its formulation. Recent studies have proposed the use of SDP and SOCP to formulate the distributed OPF [149,150]. Such formulations would also needs to take into account the power converters for the type of ANM envisaged here. For instance, a formulation similar to the one proposed in [150] could be modified to include the equations from the converters and then implement the ADMM. Nonetheless, as was discussed at the beginning of this section, the underlying assumptions in these convex relaxations would have to be assessed carefully to ensure that the operation of the LV networks is represented realistically. Having identified limitations of the traditional OPF and also in the existing alternatives, other control algorithms will be proposed in the following chapters with the aim of addressing the limitations.

## 3.7 Summary

Augmenting distribution networks with power converters and DC-links provides additional flexibility in power routing for DNOs. This flexibility is provided by the additional power flow paths and extra features, such as reactive power injections and phase-balancing. This allows a more efficient distribution of the power between feeders and a more even loading of substation transformers, creating additional headroom for connecting further LCTs. To accomplish this requires coordinated and optimised control, which is a challenge in terms of computational effort and sufficiency of network instrumentation.

This chapter has described an extension of the typical OPF formulation to become the heart of an ANM scheme for networks incorporating power converters and DC-links A flexible AC/DC OPF algorithm has been formulated, which can be used to produce optimal set-points for power converters and DC connections embedded within LV networks. It was shown that the objective function can accommodate several objectives, from losses minimisation as in section 3.4.1 and section 3.4.3 to loading equalisation between neighbouring substations in section 3.4.2. The main benefits of including power converters in distribution networks have been evaluated through the analysis of case-studies covering several network topologies, including underground cables and overhead lines, with different $X : R$ ratios and using different data availability assumptions. Also, different configurations for the power converters were used, from phase-balancers to MTDC connections. It was demonstrated that by improving the operation of the network in terms of voltage control and power routing, the power converters are capable of releasing latent capacity in the network, allowing further connections. However, the results showed that the scale of the benefits obtained depend on the characteristics of each network and that the network operator should decide between the candidate configurations based on relieving the most stringent constraint.

It was shown in both benchmark and realistic LV networks that the benefits significantly outweigh the additional losses incurred in the converters and the DC lines. It was also demonstrated that the benefits in the long-term grow as more LCTs are installed in the network. Further, given the uncertainty over the exact locations where LCTs will be connected, the converters provide additional benefits when compared to more traditional ANM schemes which use network switching. Finally, the benefits and limitations of the method chosen to solve the OPF problem were discussed and compared to alternatives found in the literature. It was shown that due to the nonlinear and non-convex characteristics of the OPF, it is difficult to scale it to large LV networks, as it becomes computationally intensive. Moreover, due to the intermittent characteristics of LCTs, calculating the optimal power distribution for every possible scenario using an OPF is not viable for real-time operation. The challenges associated with the implementation of ANM schemes in LV networks given the lack of monitoring and instrumentation at this voltage level were also discussed. It was seen that the OPF approach relies on a full-observability of the network and can thus be challenging to implement in

real-time operation given the lack of instrumentation and data communication in LV networks, as well as the short-term variability of power consumption/production of LCTs.

The following chapters will propose alternative control schemes to shift the computational burden of the OPF to an offline stage and to overcome the requirement of full observability of the network.



# Chapter 4

# Data-driven Control for ANM

As was discussed in section 3.5, an ANM scheme with power converters requires a control scheme that: 1) creates optimised set-points for the power converters based on network measurements from multiple locations and 2) updates this frequently to respond to highly variable conditions in the network. This becomes challenging if the power converters are not located at substations but instead placed downstream, in the middle or at the end of LV feeders, where the present communication and instrumentation infrastructure is poor. This chapter proposes a data-driven control strategy composed of two stages (offline and online) to respond to this challenge.

## 4.1 Regression

The three-phase AC/DC OPF algorithm introduced in Chapter 3 will be used in the first, offline, stage to determine the optimal power set-points for each PED in the network for a range of operating points. A data-set of key measurements is built by simulating several loading scenarios and applying the AC/DC OPF from section 3.3 to obtain the power converters' set points. For this, it is assumed that both the network topology and historical data are available. A multi-output regression engine then identifies relations between the set of available measurements in the network and the active and reactive power set-points for each AC/DC converter in the network. The second, online, stage of this control uses the regression model to approximate the optimal set-points for the converters during real-time operation of the network. Thus, it is postulated, the computationally-intensive stage is offline where time is less critical whereas the real-time stage can quickly respond to rapid power flow fluctuations caused by the LCTs. This facilitates the exploitation of the latent capacity within feeders and substations.

As it was discussed in section 3.1, DNOs have started installing more monitoring devices in LV network to increase their general awareness of the network status. One example of this is the OpenLV project [55]. The have also started building connectivity models for the LV network, which can be used for power flow

analysis [158, 159]. Both of these developments should underpin the deployment of ANM schemes in the coming years.

For the proposed data-driven control, it is assumed that two sets of measurements are available: the three-phase active and reactive powers leaving the LV-side of the transformer at the substation and the voltages at the AC terminals of the converters. These measurements should provide enough information on the overall loading of the substation, its level of unbalance and the extent of the voltage drop as observed at the terminals of the converter. The diagram in Fig. 4.1 illustrates where these measurements are located in the network.

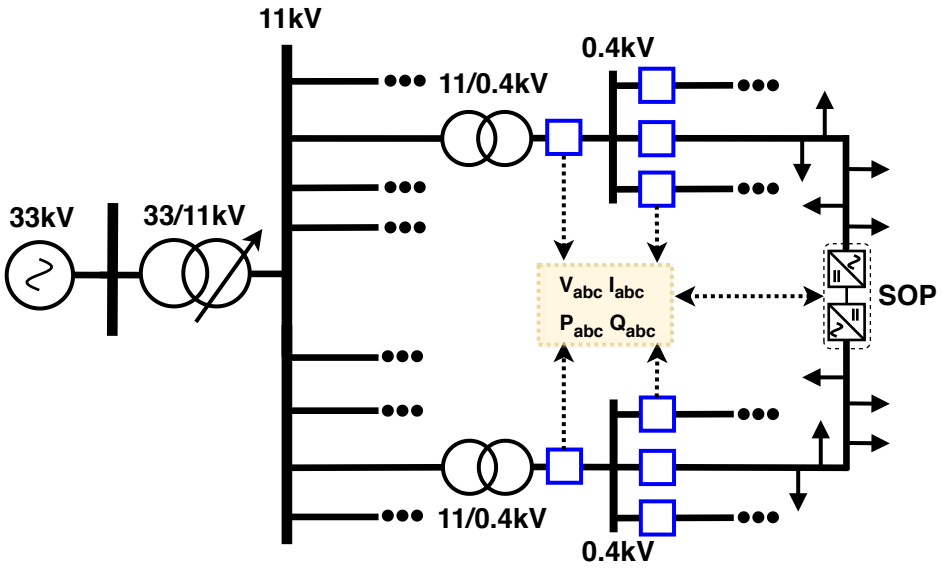



Figure 4.1: Two LV substations connected using power electronic converters, with the monitoring units shown in blue. This back-to-back connection between the converters is also referred to as SOP.

The proposed arrangement for the real-time local control assumes the existence of a communication channel between the substation, where the central controller would be located, and the converters. This communication channel will be used to transmit the power set-point instructions to the converters. However, the control will also need to take into account the potential loss of this communication channel due to equipment failure. Hence, the controller should also be able to make a reasonable estimate of the set-points for the converters based only on the voltage measurements at the terminals of the converters and peer-to-peer (P2P) communication between converters. Another assumption made is the availability of historical loading data, which can be used to train the regression algorithms. Although this might not be immediately available, the availability of data is expected to improve as the amount of monitoring devices increases. If historical data is not available, a probabilistic model of the loads and generation profiles could be used (e.g., using a Monte Carlo approach). Alternatively, state estimation algorithms in combination with smart meter data could also be used, as suggested in [160].

Several algorithms exist that can perform regression modelling. According to their characteristics, they

can be classified in parametric and non-parametric types [161]. Parametric regression models estimate the coefficients of a function $f(x)$ by selecting those parameters that minimise the error over the set of training data. The error is often formulated as a loss function, over which the minimisation is carried out. For the case of parametric methods, the functional form of the mapping between the variables and the number of parameters are fixed *a priori*. This is not the case for non-parametric models, where both the number of parameters and the shape of the model are chosen during the performance of the algorithm according to the data itself [162].

The most common and simple regression approach is linear regression, which for the case of one or more features, can be formulated as in (4.1) [161]:

$$\hat{y_i} = \hat{\beta}_0 + \sum_{j=1}^{p} X_j \hat{\beta}_j \tag{4.1}$$

where:

$\hat{y_i}$ = represents the output or target variable
$X_j$ = {1, ..., $j$} is the vector of predictors or features
$\hat{\beta}_0$ = is the intercept of the estimated model
$\hat{\beta}_j$ = {1, ..., $j$} are the coefficients estimated by the model

The least-squares method is commonly used to fit the linear model given a set of training data by minimising the residual sum of squares (RSS), calculated as in (4.2) [161].

$$RSS = \sum_{i=1}^{N} (y_j - \hat{y_i})^2 \tag{4.2}$$

[161]

In the case of multiple outputs ($y_1$, $y_2$,...,$y_n$), a linear model can be used for each of them without affecting their individual estimates [161]. Two parameters that can be used to assess the quality of the regression function obtained are the residual standard error (RSE) and the coefficient of determination $R^2$ [161]. To evaluate if a given model is over-fitting, the mean square error (MSE) for both the training and test data are calculated as in (4.3):

$$MSE = \frac{1}{n} \sum_{i=1}^{n} (y_i - \hat{y_i})^2 \tag{4.3}$$

Advanced regression techniques, such as random forests, boosted trees and artificial neural networks, are not restricted to a given linear function. Instead, the input-output relationship is represented by a combination of multiple functions acting together. When using these algorithms for regression, they are often referred to as *supervised learning*, a type of machine learning.

The random forests method, introduced in [163], is based on classification and regression trees (CART). Methods based on trees partition the feature space and fit a model in each sub-space [161]. These partitions are represented as a simple decision tree. Although function relating inputs and outputs cannot be obtained, the importance of each of the predictors or features can be extracted [164]. Three key parameters need tuning for the construction of a random forest: (1) the number of trees, (2) the depth of the trees and (3) the number of predictors sampled at each split [164]. Gradient boosted trees (GBT) is an enhancement of CART in which the gradient is boosted [161]. The boosting procedure uses the information from previous trees to grow the subsequent trees [161]. This method averages over multiple small trees, which are fitted on the residuals of the previous trees, giving it a greater granularity than simple CART methods [165]. Another approach to regression is the use of artificial neural networks. A feed-forward neural network, also known as multi-layer perceptron (MLP), is used to map input values to output values by finding the best combination of the features, producing a non-linear function that approximates the behaviour of the target variable [161, 166]. Similar to linear regression models, the neural network maps input values to output values by combining many simple functions [166]. The neural network fits the data by updating its weights or internal parameters such that the error between the prediction and the targets is minimised. Because of the recursive combination of many nonlinear functions, neural networks are normally trained using back-propagation and gradient descent optimisation, in which every hidden layer sends and receives information only from those units that are connected to it [161, 166]. In feed-forward neural networks, the information flows without feedback connections, from the input layer through the intermediate hidden layers and to the output layer [166]. Fig. 4.3 illustrates the structure of a feed-forward neural network with a single hidden layer. A neural network can learn multiple input-outputs at the same time whilst also considering the relationships between the outputs, making it suitable for multi-output regression [167]. To give more flexibility to the neural network, a non-linear activation function can be used in the hidden layer. Two common activation functions are: the hyperbolic tangent function $\tanh$ and the rectified linear unit (ReLU). The $\tanh$ will be used in the case studies in section 4.3. The transformation made on $f(x)$ by a single-layer neural network with a $\tanh$ activation function can be summarised as the operation in (4.4) and is illustrated in Fig. 4.2 [166, 168]. The $\tanh$ activation function was chosen because it can restrict the output of the neural network between $-1$ and 1. This way, the output of the power converters can be restricted to be within their rated apparent power expressed as a per-unit quantity.

$$f(x) = \tanh\left(w_0 + \sum_{n=1}^{N}(w_n \cdot x_n)\right) \tag{4.4}$$

where:

$w_0, w_n$ = represent the weights of the input and the hidden layer

$N$ = represent the number of weights of the hidden layer

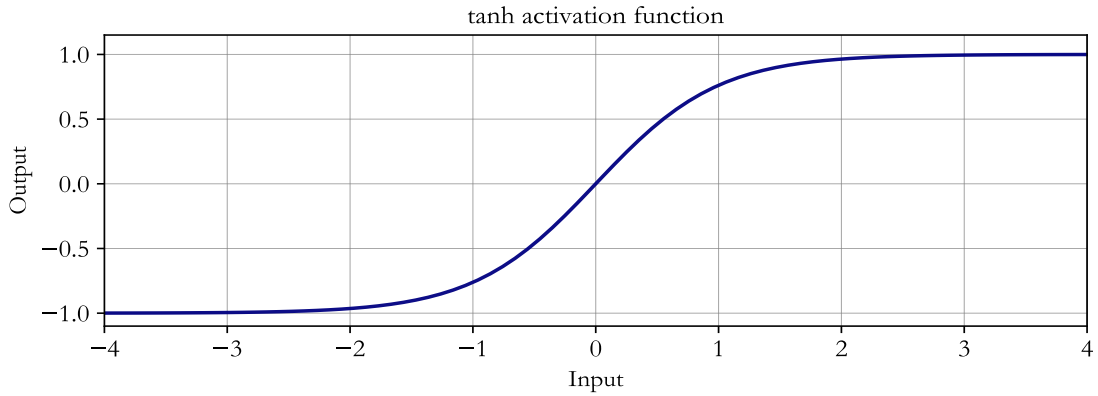



Figure 4.2: Illustration of the range of values produced by the tanh activation function.

Each of the regression models described above offers a different combination of prediction accuracy and generalisation. The models based on trees are expected to have good prediction accuracy but to generalise poorly on new data [169]. In contrast, neural networks tend to have a better performance on unseen data but are harder to tune. Regression using linear regression, random forests (RF), gradient boosted trees (GBT) and neural network will be evaluated for data-driven network control in the next section.

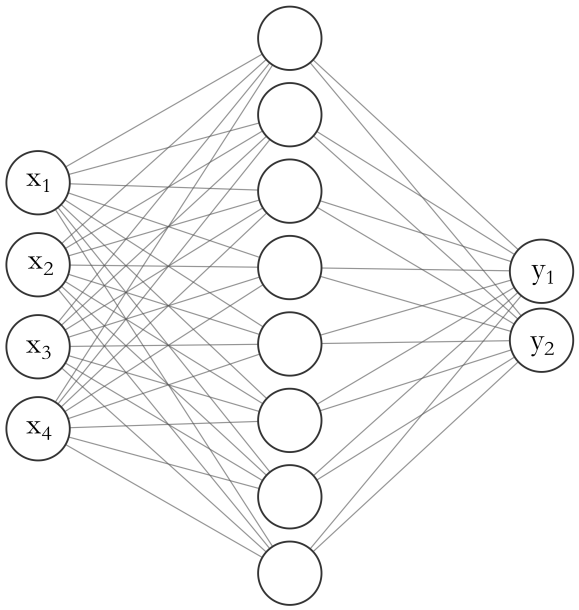



Figure 4.3: Neural network with 4 inputs, 2 outputs and a single hidden layer with 8 hidden units. This neural network has 48 weights [170].

## 4.2 Gradient-based Optimisation

The most common optimisation engine for training neural networks is based on gradient descent, or an extension of it (e.g. gradient descent with momentum, stochastic gradient descent). Gradient descent is a derivative-based (and also first-order) optimisation algorithm, which minimises a function by calculating its gradient with respect to a point and iteratively updating the function parameters in the direction of the negative gradient, towards a local minimum [166]. Given an objective or cost function $g(x)$, gradient descent can be used to find the optimal value for the variable or vector $x$ such that the cost function is minimised [166,168]. The minimum of the function can be found iteratively by taking the gradient of $g(x)$ with respect to $x$ and updating $x$ for $k$ iterations with a new value $x_{k+1}$, following the negative gradient direction, until the gradient is zero, as in (4.5) [166].

$$x_{k+1} = x - \alpha \nabla_x g(x) \tag{4.5}$$

where:

$\alpha =$ is the update step-size or learning rate
$k =$ is the number of iterations

The step-size of the updates is according to a given learning rate $\alpha$. Typically, back-propagation is used to efficiently calculate the gradient of the cost function (e.g. the error between targets and prediction) with respect to the neural network parameters or weights [166, 171]. Nonetheless, back-propagation is a general-purpose algorithm that can be used to calculate the gradient of the outputs with respect to the inputs of any system composed of differentiable functions [172]. A more detailed description of the back-propagation algorithm is given in section 5.3.1.

An immediate extension of gradient descent is to incorporate *momentum*, where the step-size varies according to the size and direction of the subsequent gradients [166]. When several gradients have the same direction, *acceleration* is applied to increase the step-size, otherwise it decreases. More advanced optimisation methods, such as AdaGrad and RMSProp, adapt the learning rate of each individual parameter as the training progresses [166]. In a further development, the authors in [173] proposed the Adaptive Moment Estimation (Adam) algorithm, combining the ideas of momentum and parameter-wise adaptive learning. Adam is widely used for training deep neural networks and thus will be used in the case studies.

It is worth noting that the optimisation algorithms used for training neural networks are sensitive to the initialisation values, due to the non-convex nature of the loss function. Therefore, the internal weights of the neural network are typically initialised to random values, whereas the biases are initialised to values close to zero [166, 174].

## 4.3 Case Studies

It was shown in the case studies in section 3.3 that minimising the power losses in a network resulted in improvements in several other features such as the voltage profile, the level of unbalance and the peak current. Building on those results, the objective set for the converters will be to minimise the power losses in the network, subject to the typical voltage constraint of remaining within ±10% of nominal voltage. The same AC/DC OPF from chapter 3 will be used. Load flow calculations will be performed in OpenDSS to generate the initial data and also to compare the performance of the regression method to the OPF [95]. Because the regression mapping is performed for each converter location individually, and it is not possible to incorporate the OPF constraints, it is possible that the combination of set-points for all of the converters would fail to observe the overall power balance that must exist in the DC system. A post-processing stage is used to overcome this deficiency and to ensure also that the converter ratings are respected. The stages of this data-driven control method are summarised in Fig. 4.4.

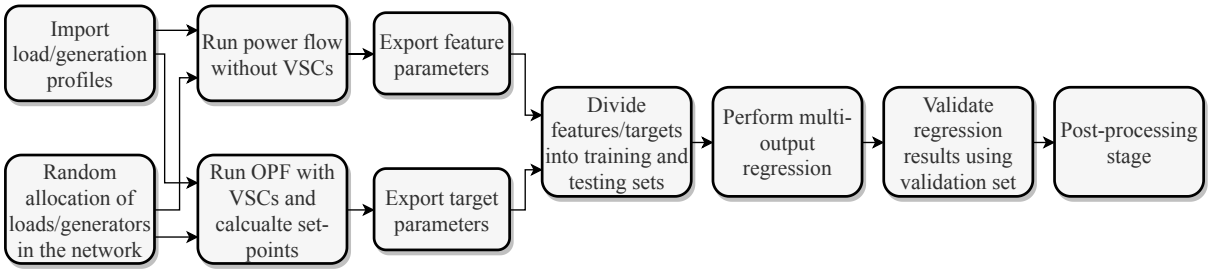



Figure 4.4: Overview of the regression training and use of the data-driven control method.

A 20-minute resolution was chosen for loads and generation profiles to capture the high variability of both. There are, therefore, 72 time-points across 24 hours of operation. The profiles were obtained by averaging a pool of individual profiles from projects in the UK [22,145]. Fig. 4.5 illustrates the various profiles.

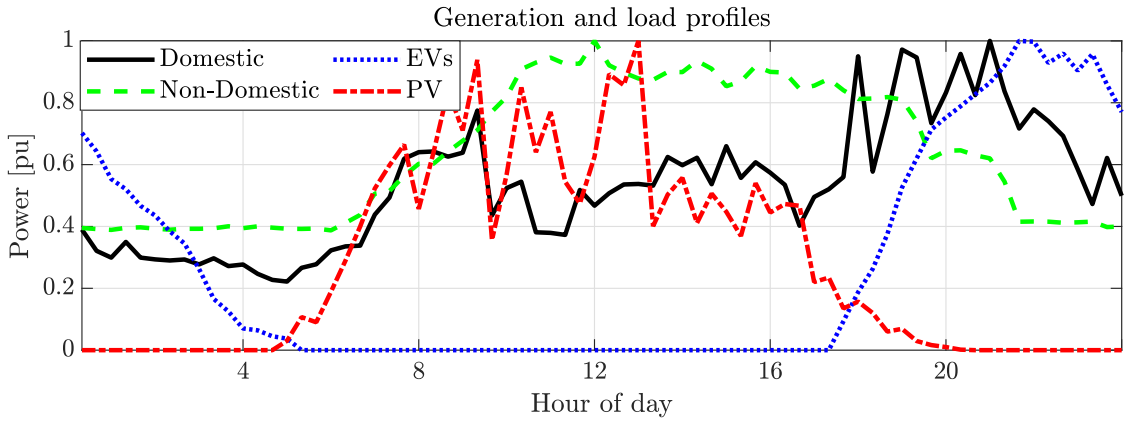



Figure 4.5: Load and generation profiles with 20-minute resolution.

To build the training dataset, several network scenarios are created. Loads and generators of the same type are assigned the same aggregated profile shown in Fig. 4.5 but each is given a different scaling factor (peak value). They are connected to a randomly-assigned bus and phase in each scenario (noting that small-scale loads and LCTs tend to be single-phase devices). The scaling factors are chosen such that the total load and generation connected to the feeders of the substation is the same across network scenarios but in each case the feeders see different operational conditions. For instance, one scenario may have loads concentrated on one particular phase: the voltage at the terminals of the converter will reflect this and the power injections determined for the converters will vary accordingly. The use of aggregate profiles with scale factors reflects the sort of information that is commonly available to a DNO. The sets of data obtained from running the OPF on these scenarios will be split into a training set (75%), a test set (25%) and a single validation example. The data will be standardised to have unit variance and a mean of zero prior to the regression, to avoid the impact of difference in units between the features [161]. The regression parameters such as number and depth of trees, hidden layers and regularisation, were tuned using a randomised search over a set of possible values. A 5-$k$ cross-validation was used with the training set to avoid under- and over-fitting [161]. The $n-k$ cross validation approach randomly splits the training data into $n$ individual sets, performs training on $n-1$ of them and uses the final set for testing. This process is repeated $n$ times, changing the order of the sets that are used for training and testing. Finally, each model was trained using the complete training set. The results of each of the regression models for the training and validation sets are presented in the case study.

### 4.3.1 IEEE 123-bus Test Feeder

The IEEE 123-bus test feeder introduced the previous chapter is used again to evaluate control performance [143]. The network diagram is reproduced here in Fig. 4.6 for ease of reference.

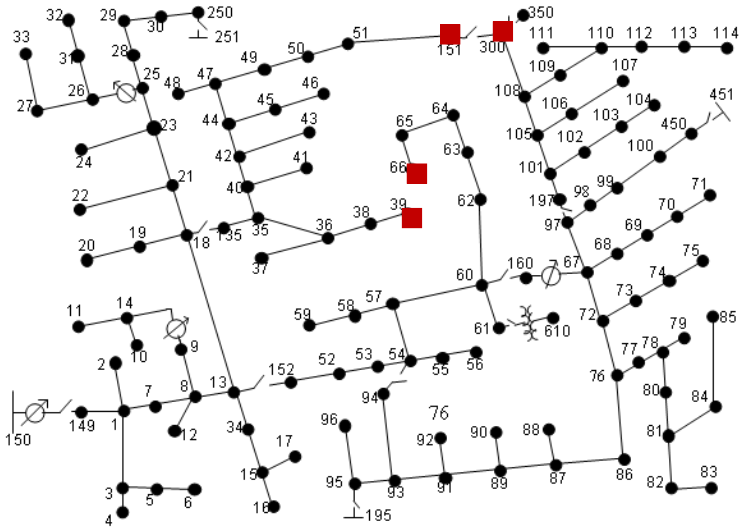



Figure 4.6: IEEE 123-bus test feeder diagram with VSCs at marked buses (39, 66, 151, 300) [143, 144].

For this evaluation, the regression model for each converter of a DC link takes as inputs the active and reactive power of its respective substation and the voltages at the terminals of both converters. The locations of the measurements (all three-phase) used to estimate the set-points are summarised in Table 4.1.

Table 4.1: Network measurements used for training the regression model.

| Measurement | Description | Target or Predictor |
|---|---|---|
| $Pvsc_a$ | Active power at VSC terminal - Phase A | Target |
| $Pvsc_b$ | Active power at VSC terminal - Phase B | Target |
| $Pvsc_c$ | Active power at VSC terminal - Phase C | Target |
| $Vvsc_a$ | AC voltage at VSC terminal - Phase A | Predictor |
| $Vvsc_b$ | AC voltage at VSC terminal - Phase B | Predictor |
| $Vvsc_c$ | AC voltage at VSC terminal - Phase C | Predictor |
| $P_{Txa}$ | Active power at Trafo LV side - Phase A | Predictor |
| $P_{Txb}$ | Active power at Trafo LV side - Phase B | Predictor |
| $P_{Txc}$ | Active power at Trafo LV side - Phase C | Predictor |
| $Ploss_{Tx}$ | Active power loss of Trafo | Predictor |

The converters can communicate their respective local voltage measurement with each other. The three active and reactive powers for each phase of each converter will be used as target variables. The results of the training and validation of the regression models are presented in Table 4.2. It can be seen that despite achieving good prediction results on the training set, the tree-based models, RF and GBT, do not generalise well to the validation set and give an accuracy $R^2$ of only 0.218 and 0.259, respectively, and are not considered further.

Table 4.2: Training and validation results for the four regression methods.

| Regression Model | $R^2$ of the training set | MSE of the training set | $R^2$ of the validation set | MSE of the validation set |
|---|---|---|---|---|
| Linear | 0.908 | 0.169 | 0.673 | $1.04\times10^{-4}$ |
| Neural network | 0.935 | 0.137 | 0.757 | $1.04\times10^{-4}$ |
| RF | 0.929 | $4.46\times10^{-5}$ | 0.218 | $1.63\times10^{-4}$ |
| GBT | 0.916 | $5.39\times10^{-5}$ | 0.259 | $1.78\times10^{-4}$ |

The best balance between prediction and generalisation was offered by the neural network model, with a good $R^2$ score of 0.757 on the validation set. The neural network parameters obtained from the randomised search for this case were: 3 hidden layers, each with 110 weights, using a $\tanh$ activation function and optimised using the Adam optimiser with a learning rate of 0.6. The neural network model was selected for further validation by comparing the converter set-points it produces against the OPF. The histogram in Fig. 4.7 shows the error between the set-points produced by the regression model and the OPF for both the converters. The mean error is close to zero with only a few occurrences of an error greater than $\pm$0.01 MW/Mvar, which is equivalent to 5% of the converter rating.

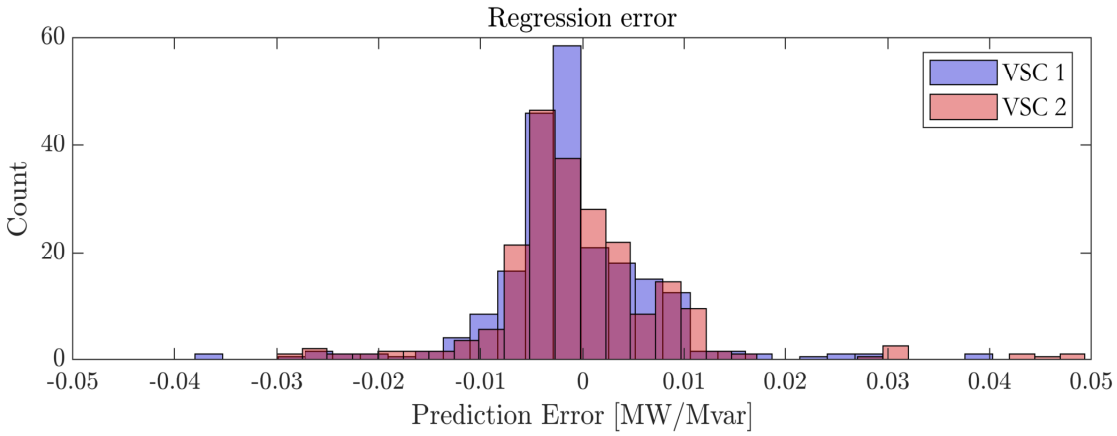



Figure 4.7: Histogram of error of neural network regression model with respect to the OPF for the IEEE 123-bus test feeder.

Fig. 4.8 compares the active power injections of each phase of both the converters (one after the other) over 24 hours of one of the validation cases. Some minor offsets are observed between the regression model and OPF results but the regression model successfully captures the main features of the OPF results. Similar good agreement was found for reactive power injections and the combined real and reactive power injections do not exceed the rating of the converters.

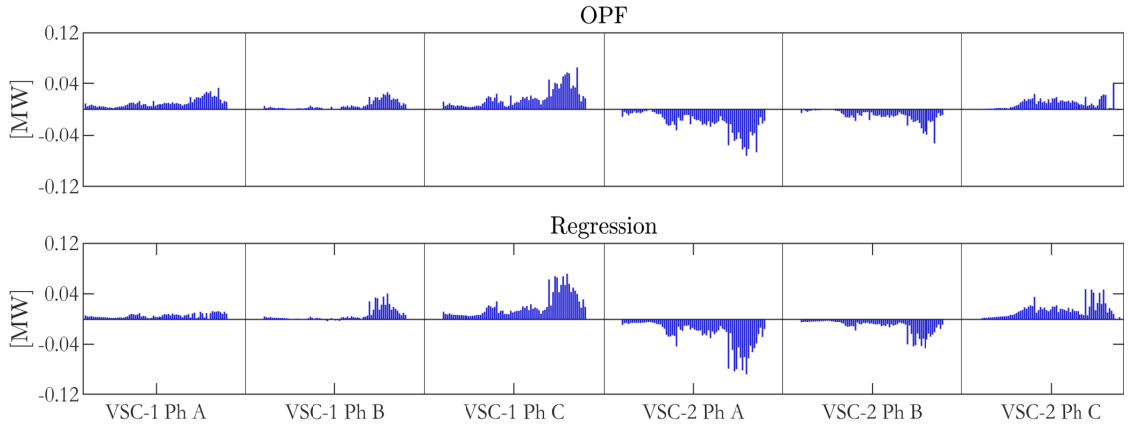



Figure 4.8: Comparison of active power set-points from regression model and OPF of each converter phase (set out sequentially for the 24hr period).

A power flow was performed using the set-points obtained from the regression model for one of the validation cases and the currents observed at the head of the feeder for the OPF and the regression model are shown in Fig. 4.9. Both methods reduced peak phase current and reduced the neutral current compared to a base case where DC connections were not used. There is also a decline in the voltage unbalance (VUF) measured at the converters' terminals, indicating that unbalance is improved along the feeder. The regression method performs broadly as well as the OPF.

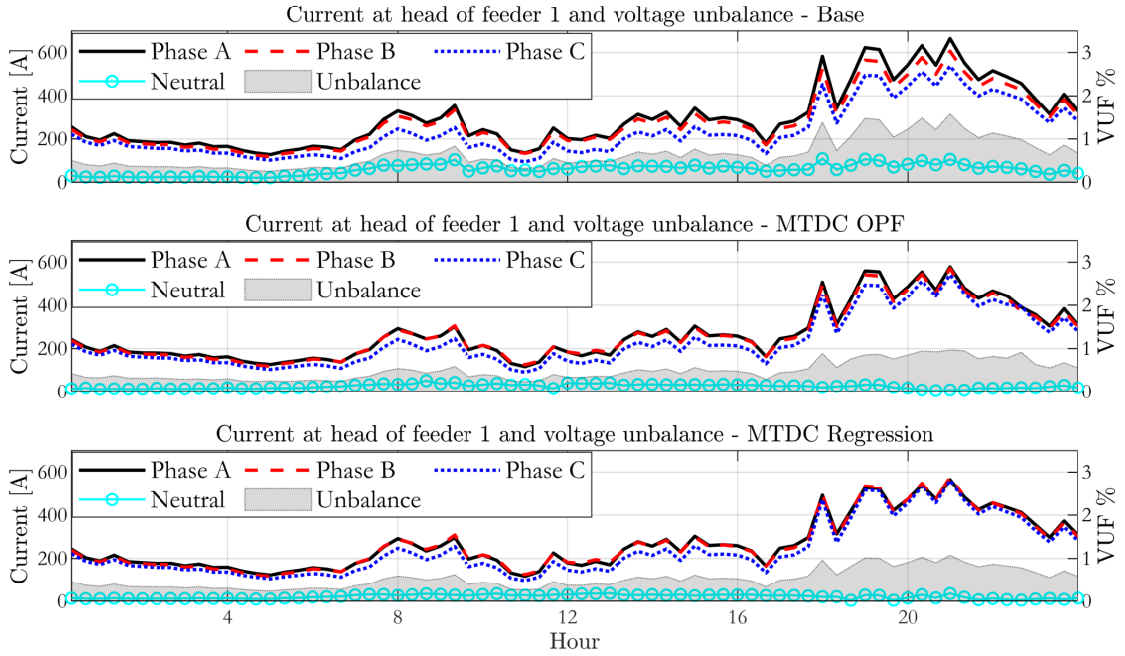



Figure 4.9: Comparison of the current at head of feeder and voltage unbalance factor for the IEEE 123-bus test feeder.

Table 4.3 shows that other parameters such as minimum voltage and power losses are improved considerably by the regression method, as well as the OPF. One node in the network was originally operating below its voltage limit of 0.9 pu and both the regression method and the OPF eliminate this limit violation. In the base case, the transformer loading was at 0.913 with little capacity left for new connections. However, the case including DC connections frees up 0.05 pu of capacity under both control methods, as a result of the lower reactive power requirements. This is also reflected in a reduction in peak current at the feeder head as seen in Fig. 4.9.

Table 4.3: Comparison of regression method performance against OPF and base case for IEEE 123-bus test feeder.

| Case | Voltage min [pu] | Losses in network [MWh] | Losses in VSCs [MWh] | Transformer Loading max [%] | Current max [A] |
|---|---|---|---|---|---|
| Base | 0.875 | 1.395 | - | 91.3 | 664.5 |
| OPF | 0.905 | 1.277 | 0.249 | 85.7 | 578.2 |
| Regression | 0.910 | 1.289 | 0.253 | 86.1 | 576.5 |

Not only do these results verify the effectiveness of the converters control using the regression model but also that the control can be based on measurements at only the substation and the converter stations.

### 4.3.2 European LV Network

The realistic LV network from section 3.4.3 is used here, together with another similar feeder adapted from [175], to further demonstrate the capabilities of the regression method proposed. In this case, the loads in feeder 1 have residential profile, with a total peak load of 0.191 MW, whereas those in feeder 2 have a non-residential profile, with a total peak load of 0.286 MW, as illustrated in Fig. 4.5. Consequently, the two substations will have different loading patterns. Both types of loads are assumed to have a power factor of 0.95. In addition, a set of 45 EVs, each with a 3.3 kW demand, are randomly assigned across the two feeders and 30 PV generators with a 6 kW capacity are allocated to feeder 1.

The two networks can be connected via a pairs of power converters at the points where they are physically close together, as illustrated in Fig. 4.10. The rated power for the three-phase converters was set at 120 kW. However, reactive power will not be used because the LV lines are mostly resistive (the $X : R$ ratio is low) and it is expected that reactive compensation will not be effective for improving the voltage profile. It is better to reserve the converter ratings for active power only.

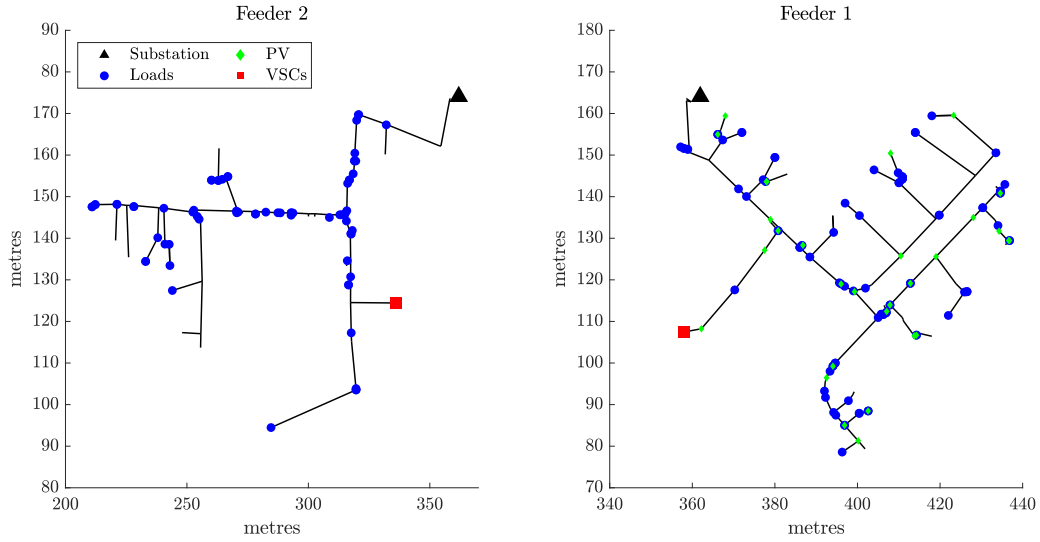



Figure 4.10: European LV network diagram with VSCs at marked buses [175].

The regression model (and OPF) required remote measurements from the substation to operate the DC-link and so vulnerability to loss of communications is a concern. The proposal is to train an alternative regression model based on neural networks that does not rely on remote measurements from the substations and uses only local measurements. This regression model can be substituted whenever communication loss is detected. The results of this alternative are also evaluated and compared to assess if the performance is adequate.

On average it took approximately 3-5 minutes for the AC/DC OPF to solve each of the 20-minute time-step of a 24 hour period. The number of scenarios used for training was limited to 20 given the time taken per solution and as the performance during validation was found to be satisfactory. A summary of the results

for each regression model is presented in Table 4.4.

Table 4.4: Training and validation results for the EULV network.

| **Regression Model** | $R^2$ **of the training set** | **MSE of the training set** | $R^2$ **of the validation set** | **MSE of the validation set** |
|---|---|---|---|---|
| Linear | 0.810 | 0.181 | 0.524 | $2.65\times10^{-5}$ |
| Neural network | 0.854 | 0.139 | 0.681 | $1.96\times10^{-5}$ |
| RF | 0.846 | $1.64\times10^{-5}$ | 0.568 | $2.23\times10^{-5}$ |
| GBT | 0.841 | $1.68\times10^{-5}$ | 0.542 | $2.44\times10^{-5}$ |

In the EULV case, the regression models is slightly less accurate than it was for the IEEE 123-bus test feeder. The neural network has the highest validation $R^2$ of 0.681. The reason for this might be that as the number of nodes in the network increases, the complexity of the model must increase as well. The same neural network parameters were used, as in the IEEE 123-bus test feeder case.

Fig. 4.11 shows the distribution of error for the validation scenarios. Again, the error is seen to be slightly larger than in the IEEE 123-bus case. The error is still centred around zero with only a few occasions in which the error is greater than $\pm5$ kW, 12.5% of the converter rating. The error for the case with partial communication is also depicted in Fig. 4.11 where it can be seen that the impact of loss of communications is small.

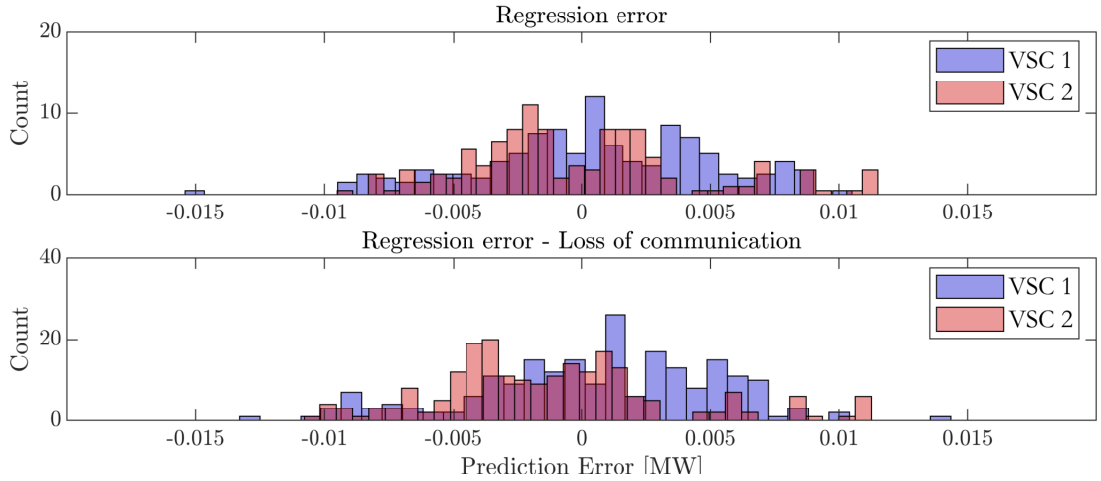



Figure 4.11: Comparison of active power set-points from regression model and OPF of each converter phase (set out sequentially for the 24hr period).

The power set-points of the converters are compared for OPF and regression model in Fig. 4.12. The regression model follows the trend of the OPF closely, however it does present some minor differences. The impact of these errors is explored in Table 4.5 and Table 4.6, which record the network performance for base-case, OPF and the regression model. The impact of a loss of communications is also included in these tables. Given the small differences with the main regression, operation during an interruption to communications is likely to be acceptable.

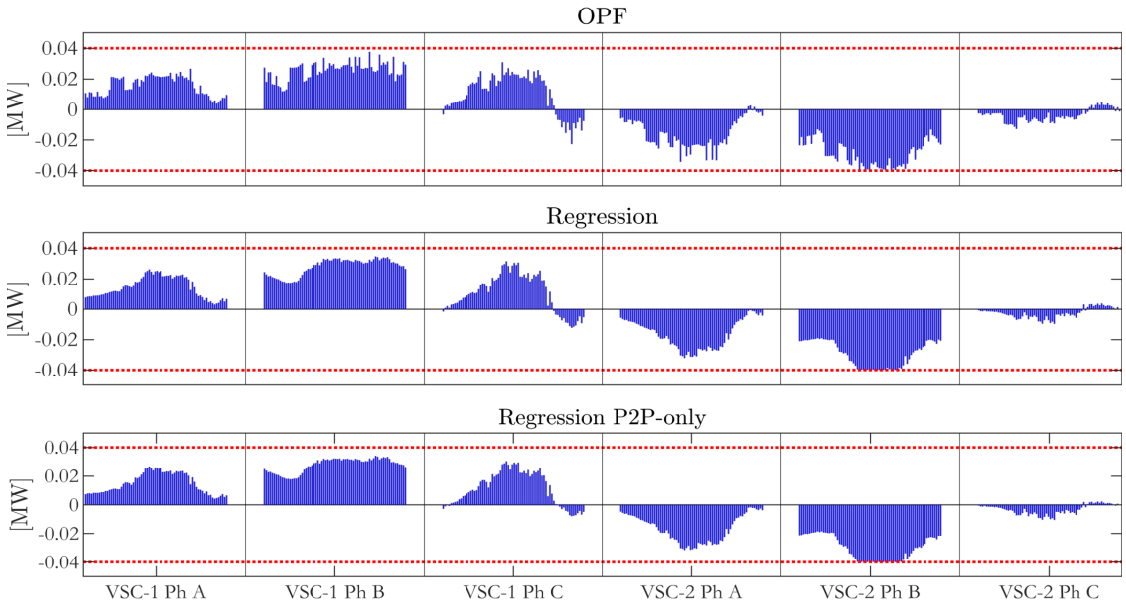



Figure 4.12: Comparison of active power set-points from regression model and OPF of each converter phase (set out sequentially for the 24hr period).

Table 4.5: Comparison of regression model performance in EULV Network.

| Case | Voltage min [pu] | Losses VSCs [MWh] | Total losses [MWh] |
|---|---|---|---|
| Base | 0.868 | - | 0.372 |
| OPF | 0.910 | 0.091 | 0.381 |
| Regression | 0.905 | 0.094 | 0.385 |
| Regression No-Comms | 0.905 | 0.095 | 0.386 |

Table 4.6: Comparison of regression model performance in EULV Network - Feeder 1 and 2.

| Case | Losses F1 [MWh] | Losses F2 [MWh] | Transformer F1 Loading max [%] | Transformer F2 Loading max [%] | Current F1 max [A] | Current F2 max [A] |
|---|---|---|---|---|---|---|
| Base | 0.088 | 0.284 | 95.3 | 103.9 | 562.8 | 589.1 |
| OPF | 0.104 | 0.186 | 103.3 | 79.5 | 490.8 | 404.3 |
| Regression | 0.107 | 0.184 | 105.6 | 78.3 | 515.0 | 411.5 |
| Regression No-Comms | 0.108 | 0.183 | 107.2 | 77.2 | 529.7 | 411.2 |

The active power losses are slightly higher for feeder 1 when the DC-link is used (both OPF and regression compared to base case) because of the power losses of the converters. Nonetheless, the losses of the converters represent a small portion of the total losses as it was shown in Table 4.5.

As shown in Table 4.6, the active power losses are reduced significantly for feeder 2, especially during the hours of peak solar generation, because of the complimentary profiles of the two feeders. Feeder 2 has a peak load during daylight hours so the export of PV generation from feeder 1 to feeder 2 lowers reactive power consumption in the transformers and reduces their loading. As an example of the improved phase-balance and the reduction in the current magnitude at the head of the feeders, the currents at feeder 2 are depicted in Fig. 4.13.

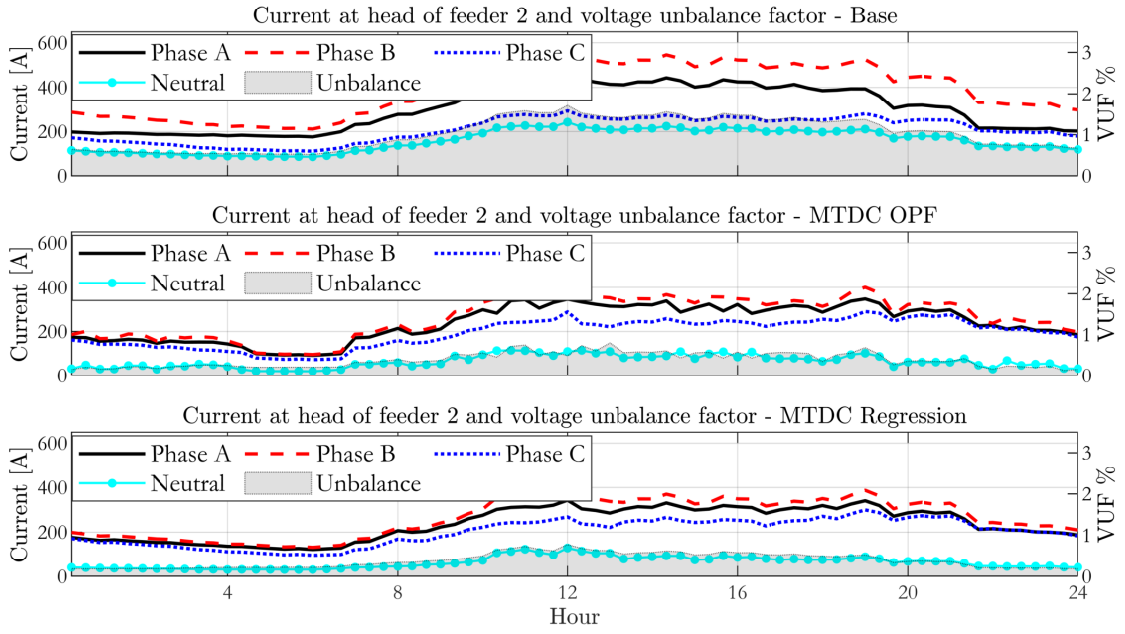



Figure 4.13: Current at head of feeder 2 for base case, OPF and regression model.

## 4.4 Discussion

The results from both case studies demonstrate that the regression model using neural network performs relatively well and closely approximates the OPF with only a few instances of small divergence. An error of approximately 5% for the IEEE 123-bus case and 12% for the European LV case. Notwithstanding these promising results, it is important to note that the regression model is not suitable for extrapolation to cases outside the range of the training data. If an unusual situation were observed, the set-points provided by the regression model could be far from optimal. For instance, if the network configuration changes due to switching, or if the overall generation increases, the regression model would no longer be valid.

Because of this, it would be sensible to detect conditions outside the training range and suspend power

exchange between the converters. This could be followed by running the OPF for the specific conditions and expanding the training so that future occurrences can be accommodated.

It's worth highlighting that the regression model represents an open-loop approximation to the optimal control and it does not update its output according to feedback from the network model. This could lead to inappropriate behaviour if abnormal situations occur, such as faults. Further studies will be needed to explore these conditions.

## 4.5 Summary

This chapter has proposed a data-driven control strategy to overcome some of the difficulties in computational effort and sufficiency of network instrumentation that prevent the use of ANM based on a real-time OPF. The control comprises two-stages: an offline stage where the OPF is used to generate optimal converter set-points and an online stage where the regression model is used to approximate these in real-time. This combines the advantages of a high-level optimisation with the efficiency of a data-driven technique using only a reduced set of measurements and with a low communication requirements. A review of the literature on linear and nonlinear regression and its application to power systems was carried out in this chapter. From this, the regression algorithms that might be used to accurately map the interaction between network variables and optimal power set-points were identified.

It was demonstrated through case studies that despite their acceptable accuracy, regression models are not suitable for extrapolation to cases outside the range of the training data. Moreover, the regression model represents an open-loop control, which does not update its output according to feedback from the network. This could lead to inappropriate behaviour if abnormal situations occur, such as line outages or network switching. With this in mind, the next chapter will propose alternative control methods based on a more sophisticated form of machine learning, which allows constant feedback from the system to tune the response of the control scheme and enable a form of generalisation.



# Chapter 5

# Reinforcement Learning Control for ANM

Some of the results in this chapter are taken from the following paper:

- J. Perez-Olvera, T. C. Green, and A. Junyent-Ferre, "Self-learning Control for Active Network Management," in *2021 IEEE Madrid PowerTech*, Madrid, Spain, 2021 (Accepted)

Part of the work on the model-based approach reported in this chapter was carried out in collaboration with Dr Michael Fairbank, a lecturer at the University of Essex. Dr Fairbank had developed similar optimisation algorithms for, amongst other things, motor drives. The incorporation of a distribution network power flow into the model-based optimisation framework was conducted with his help and advice. The production of the results for the active network management study and their interpretation are the work of the author alone.

## 5.1 Reinforcement Learning

Rather than trying to map measurements to optimal settings obtained from traditional optimisation using regression, the optimisation process can be formulated using an entirely different framework. Prompted by the success of reinforcement learning applied to high-dimensional control problems [176, 177] and dynamic optimisation [178] this chapter proposes using reinforcement learning as an alternative control approach. Reinforcement learning has not so far been suggested as a framework to solve the OPF problem, nor has it been applied to scheduling the set-points of power converters for ANM. In this chapter, a study of the application of reinforcement learning to ANM will be reported using two case studies.

Assumptions of full visibility that underpin the application of an OPF are not realistic in practice. Thus, the reinforcement learning methods will be assessed for their ability to work with only partial visibility of the state of the distribution network, as discussed in sections 3.1 and 3.5. Furthermore, unlike regression algorithms, reinforcement learning methods receive constant feedback from a model of the system, which

shapes the the control policy and offers the ability to generalise [179, 180]. The type of control proposed in this chapter can be also be described as adaptive because it start with an initial baseline performance that improves based on experience.

The goal of a reinforcement learning algorithm is to train an agent to find a control policy that minimises a given objective function. This can be done through direct interaction with the equations describing a system (model-based) or, alternatively, through interaction with measurements from an actual system or a *black-box* representation of the system (model-free) [181]. In the context of power systems, recent applications of reinforcement learning have been limited to electricity market design [182], demand response problems [183], energy management systems for microgrids [184] and, recently, energy management for buildings [185].

The study in [182] uses Q-learning, enhanced with *virtual* experience, to develop a dynamic pricing systems for the energy supplier and for the consumers in a microgrid. This *virtual* experience allows the agents to take advantage of multiple state-actions pairs, as well as information on the state transition probability, thus speeding up the agents' training. The main goal for the supplier is to price the energy based on the total demand observed in the microgrid. For the consumers, the authors adopt a multi-agent approach in which each customer is considered as a learning agent that adapts their consumption based on the price observed . As a model-free implementation, the system works without *a priori* information on the microgrid dynamics. Although the results presented are positive, it is not clear how the algorithm proposed would scale to realistic microgrids, if all the state-action pairs need to be stored in memory, given that it does not use any form of function approximation.

The demand response problem for residential thermal loads is addressed in [183] through the use of an extended version of fitted Q-iteration (introduced in [186]). The authors extended fitted Q-iteration by incorporating information from expert knowledge and by providing forecasts of some external variables, which directly affect the system, to the learning agent. The results of the proposed approach based on RL are compared to an optimal controller with full access to the model parameters. Although a performance difference of more than 20% was reported, between the optimisation and the RL approach, the learning agent managed to reduce the total cost of electricity consumption.

The authors in [184] implemented a policy iteration method based on approximations of the Q-function to optimise the energy in a microgrid. The Q-functions are approximated at every iteration using extremely randomized trees (an extension of random forests, described as in section 4.1). Nonetheless, the algorithm proposed was shown to be computationally demanding, given its model-free nature, taking more than 6 hours to train.

In the field of optimal control and learning agents, there are two main frameworks, which often intersect but have different backgrounds. These two frameworks share the same objective: to train an agent to make

decisions that maximise a given objective function [187]. On the one hand, there are the techniques based on Reinforcement Learning [181] and Markov Decision Processes (MDPs). On the other hand, a framework for training a learning agent emerges from the design of optimal controls for dynamic systems and the concept of dynamic programming [187–189].

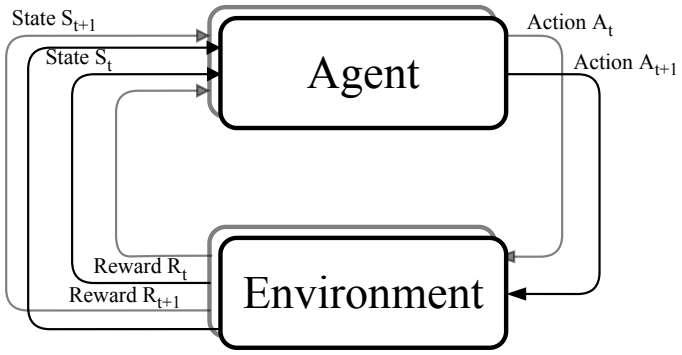



Figure 5.1: Agent-environment interaction over multiple time-steps in a Markov process [181].

Dynamic programming, as defined by Bellman in [180], attempts to solve problems involving a sequence of optimal decisions and the effects of these decisions at a later stage (i.e. long-term effects). This sequence of decisions can be seen as a policy, or as a set of steps that result in optimal behaviour of a given system, for every state of the system [180, 190]. Calculating the optimal control policy over a trajectory for a nonlinear stochastic system can be done by traditional dynamic programming. However, this requires storing in memory both the optimal and suboptimal trajectories, making this intractable in practice due to the large number of possible states and their respective policy [191]. Approximate, or Adaptive Dynamic Programming (ADP) uses function approximators to overcome the *curse of dimensionality* associated with classic dynamic programming [172, 180, 189, 190, 192]. The term *function approximator* is used here to refer to algorithms (e.g. neural networks) that can approximate any function given enough parameters (e.g., a given number of hidden layers and neurons) [193].

## 5.2 Model-free Reinforcement Learning for ANM

Model-free reinforcement learning is a departure from typical control strategies in that it is able to realise near-optimal policies without an explicit model of the system to be controlled simply by interacting with a *black-box* representation of it [162, 181]. The term *black-box* is used here to refer to a model or software simulation of a system, which can output the state of such system given a set of inputs, for example a power flow engine. A popular form of model-free reinforcement learning is Q-learning. Initially proposed in [194] as a look-up table algorithm that employed the Markov Decision Process (MDP) framework, Q-learning has recently been extended for use in combination with neural networks as a control approach [176].

In the example of an ANM scheme, the optimal distribution of the power flows in a LV network can be framed as a sequential decision-making process and formulated as a MDP [181]. Thus, it can be represented by a tuple $(S, A, T, R)$, where: $S$ is a finite set of states $\forall\ s \in S$, $A$ is a finite set of actions the agent can take $\forall\ a \in A$, $T : S \times A \times S \rightarrow [0, 1]$ represents the transition probability from state $s_t$ to state $s_{t+1}$ when the agent takes action $a_t$, $R : S \times A \times S \rightarrow \mathbb{R}$ represents the reward function $r(s_t, a_t)$, where $r_t$ is the immediate reward obtained when taking action $a$ in state $s_t$ and progressing to state $s_{t+1}$. The mapping from environment states to the actions is known as a policy [181]. A policy-making agent can learn an optimal control policy, (set-points for power converters) by interacting with the environment (network) and by the reward (feedback on power losses or voltage accuracy) received. An illustration of this approach used for ANM is presented in Fig. 5.2. One of the advantages of the model-free approach is that a full model of the network with the individual loads is not required, aggregated load measurements, such as those from the OpenLV project [55], can be used.

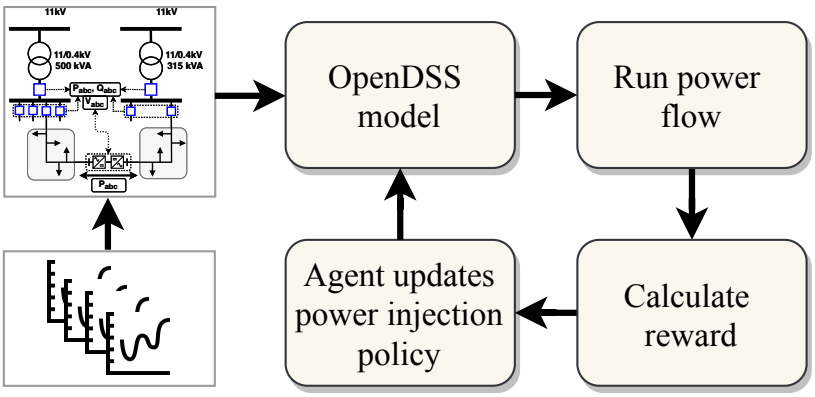



Figure 5.2: Overview of the model-free training environment for an ANM. The substation transformer and outgoing feeders are modelled in OpenDSS. The aggregated load supplied by the feeders can be obtained from monitoring devices at the LV substation (e.g., monitoring devices installed in the OpenLV project [195]). The interactions between the agent and the environment are performed via power flow simulations and are used to train the neural network to make optimal decisions.

In the example to be used for this study, the agents controlling the power converters are trained to equalise the loading of two LV substations. The expectation is that equalisation of loading will reduce the overall power losses and also improve the network utilisation and the hosting capacity. The state $s_t$ will comprise a series of network measurements, which are available to the DNO via monitoring devices. These measurements describe the loading of the substations and the aggregated load supplied by each of the outgoing feeders, as well as the resulting power losses. Example measurements are the power leaving the LV side of transformer $k$ at phase $n$ ($Tx_kP_n$) and the voltage measured at the AC terminals of the converter $k$ at phase $n$ ($VSC_kV_n$). The state at every time-step is set out in (5.1) and the measurements to be used are set out in Table 5.1:

$$s_t = (Tx_k P_n, Tx_k S_{pu}, Tx_k P_{loss}, VSC_k V_k) \tag{5.1}$$

The action $a_t$ represents the three-phase power injection of both converters $VSC_k P_n$, which translates into a certain amount of power flowing from one substation into another.

The objective of the reinforcement learning problem is to learn the optimal control policy by maximising a reward signal as described in (5.2) [194].

$$Q(s,a) = r + \gamma \max_{a_{t+1}} Q(s_{t+1}, a_{t+1}) \tag{5.2}$$

where $\gamma$ represents a discount-rate parameter, which serves to balance the value of immediate rewards versus future rewards. The function $Q(s_{t+1}, a_{t+1})$ or Q value represents the value of taking action $a_{t+1}$ in state $s_{t+1}$. In this context, an optimal control policy represents any policy that maximises the cumulative reward over the episodes or epoch [196]. This can be achieved by first estimating the values of each state-action pair for a policy $\pi$, as in (5.3) [181, 194],

$$Q^{\pi}(s,a) = \mathbb{E}_{\pi}[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} \mid S_t = s, A_t = a] \tag{5.3}$$

where $\mathbb{E}_{\pi}$ denotes the expected value when the agent follows policy $\pi$. The Q values are then maximised using (5.4) and the action with the highest Q value is selected by the agent.

$$Q^{\pi}_{*}(s,a) = \max_{\pi} Q^{\pi}(s,a) \tag{5.4}$$

In Q-learning, the algorithm starts with a random Q-function and iteratively updates it using the observed transitions, as in (5.5) [181].

$$Q(s,a) \leftarrow Q(s,a) + \alpha[r + \gamma \max_{a} Q(s_{t+1}, a) - Q(s,a)] \tag{5.5}$$

Q-learning is a model-free, off-policy algorithm [181, 194], in which an explicit model of the environment dynamics is not needed, and the optimal policy is obtained by following an exploration policy, such as $\epsilon$-greedy. To make sure the agent explores the state-action space adequately, an $\epsilon$-greedy policy is used to select the action $a_t$ at each iteration, as in (5.6).

$$a_t = \begin{cases} \underset{a}{\operatorname{argmax}}\, Q(s_t, a_t) & \text{with probability } 1-\epsilon_t \\ \text{random}(a_t) & \text{with probability } \epsilon_t \end{cases} \tag{5.6}$$

The explore-exploit ratio $\epsilon_t$ starts at a value of 1 (i.e. choosing random actions) and decreases at every iteration, following the schedule set by the decay ratio $\tau$, as in (5.7). By favouring exploration at the beginning of the training stage and slowly switching to exploitation, the agent is able to sample the solution space whilst converging to the optimal control policy.

$$\epsilon_{t+1} = \epsilon_t \cdot \tau \tag{5.7}$$

Following the $\epsilon$-greedy policy, the agent explores each episode ($\epsilon > 0.5$) for $k$ iterations before turning to exploitation ($\epsilon < 0.5$) and reaching its final value ($\epsilon_{min} = 0.01$). A detailed description of the explore-exploit dilemma and its relationship to learning agents is given in [181].

### 5.2.1 Deep Q Learning Using Neural Networks

As it was described in section 5.1, estimating and storing the optimal values for all the state-action pairs is intractable due to the large state space of the OPF problem. This limitation can be overcome with the use of function approximators that estimate the Q function [197]. Deep neural networks can be used for this purpose, as it has been shown that they can approximate arbitrary non-linear functions, given enough parameters and hidden layers [193]. Hence, a deep neural network with $h$ hidden layers and weights $\theta$ is used to approximate the Q values. The fully-connected neural network approximates the value function Q, as in (5.8), and is known as a deep Q network (DQN) [176].

$$Q(s_t, a_t|\theta) \approx Q^{\pi}(s_t, a_t) \tag{5.8}$$

The input to the neural network is the state of the LV network, which can be represented by a vector of measurements available to the agent (power converters). For instance, these measurements can include the three-phase active power supplied by each transformer and the voltage at the terminals of each converter. Table 5.1 summarises the available measurements included in the input vector to the NN, according to the data available from the OpenLV project [55]. To avoid instability when training the NN, all the individual input measurements are normalised to be between 0 and 1, using their respective per unit bases. It is assumed that there is a communication channel between the substation and the converters, so that these measurements can be transmitted to the converters.

The output of the neural network corresponds to the approximated Q-value of each action, based on the current state. The actions represent the converters' power injections, discretised into steps of variable size (small, medium and large) that get accumulated as the agent iterates. Without this, an optimisation would be needed at every step, given a continuous action space. A sample of the actions is presented in Table 5.2.

Table 5.1: Summary of LV network measurements.

| Measurement | Description |
|---|---|
| $Tx_1P_{abc}$ | Three-phase power at LV side of transformer 1 |
| $Tx_2P_{abc}$ | Three-phase power at LV side of transformer 2 |
| $Tx_1S_{pu}$ | Loading of transformer 1 |
| $Tx_2S_{pu}$ | Loading of transformer 2 |
| $Tx_{12}S_{pu}$ | Difference between transformers loading |
| $Tx_1P_{loss}$ | Power losses of transformer 1 |
| $Tx_2P_{loss}$ | Power losses of transformer 2 |
| $VSC_1V_{abc}$ | Three-phase voltage at terminals of converter 1 |
| $VSC_2V_{abc}$ | Three-phase voltage at terminals of converter 2 |

A smaller (0.025 kW) and a larger (0.3 kW) step size are also used but for brevity they are not included in Table 5.2. An overview of the proposed control is shown in Fig. 5.3.

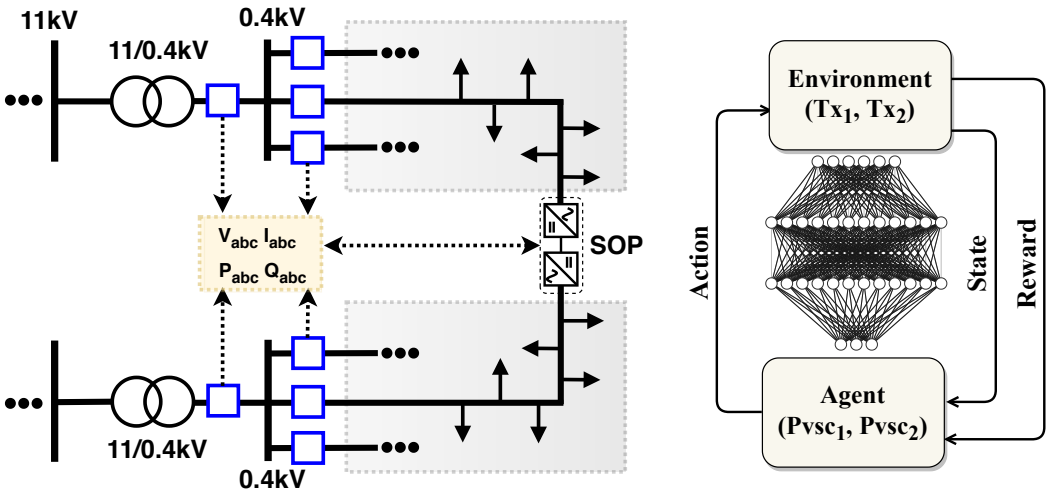



Figure 5.3: Overview of the proposed control based on deep Q learning. The LV network, composed of two substations and monitoring devices at the LV side of the transformer and at the head of the outgoing feeders, represent the environment. The power converters represent the learning agent and their power injections are the actions. The neural network approximates the value of the state-action pairs.

At every iteration the agent chooses a single action from Table 5.2, which is then added in a cumulative manner and assigned to both converters with their respective power flow direction. During training, the agent stores the observed transition $(s_t, a_t, s_{t+1}, r_t)$ at each step in a cyclic memory buffer (also known as experience replay) comprising the most recent observations [176], similar to the *virtual* experience adopted in [182]. It then uses a random batch of transitions, taken from this memory, to iteratively update the weights $\theta$ of the neural network by minimising the loss in (5.9).

$$L = (Q(s_t, a_t|\theta) - y_t)^2 \tag{5.9}$$

where $y_t$ is the target and is calculated for every transition following (5.10)

$$y_t = \begin{cases} r_t & s_{t+1} \text{ is terminal} \\ r_t + \gamma \max_{a_t} Q(s_{t+1}, a_t|\theta') & s_{t+1} \text{ is not terminal} \end{cases} \tag{5.10}$$

The memory buffer allows the agent to train on past observations [196], making a more efficient use of the data generated by the agent-environment interaction and reducing the total number of iterations needed to converge. It also makes training possible by breaking the correlation between consecutive observations [176].

Table 5.2: Discretized action space sample.

| Action | Phase A [kW] | Phase B [kW] | Phase C [kW] |
|---|---|---|---|
| 1 | 0.1 | 0 | 0 |
| 2 | 0 | 0.1 | 0 |
| 3 | 0 | 0 | 0.1 |

In contrast to supervised learning (i.e. regression), where the targets are fixed at the beginning of the training stage, in reinforcement learning (RL) the targets are constantly changing [176], as the agent explores the solution space and updates the neural network weights. Also, training/testing labels on what behaviour is expected are not provided to the agent beforehand, as these are not normally available. Due to these characteristics and to ensure stability whilst training the agent, a second neural network is used to keep a lagged copy of the targets [176]. This secondary neural network is used to calculate the Q values of the state-action pairs, which make up the targets used for training. Hence, the weights $\theta'$ of this neural network are updated less frequently, only once every $C$ number of iterations, in order to keep the targets fixed. It is worth noting that prior to training, the weights $\theta$ of the neural network are initialised using a normalised random initialisation [174].

### 5.2.2 Episodes and Termination Statements

The self-learning agent interacts with and receives feedback from the LV network model at discrete steps throughout the day. To allow the agent to learn an accurate representation, whilst keeping a reasonable training time, the interaction occurs at 5-minute time-steps, for a total of 288 time-steps. After every interaction, the agent gets the subsequent state of the network and feedback on the action taken in the form of a reward. This reward signal is a scalar value that represents the agent's goal. In this case, given the results in previous sections, the reward signal chosen is a function of both the loading difference between the transformers at each LV substation and the power losses of the transformers. Thus, if the action by the converters results in improved loading equalisation or reduced power losses, then a reward is received, otherwise a penalty is given to the agent. However, the loading equalisation takes a larger reward/penalty than the power losses, as it is the primary goal for the algorithm.

For the agent, each 5-minute step is considered as an episode in which it can reach the objective in a certain number of iterations. The agent interacts with the environment for as many iterations as needed until it reaches a termination state. A 5-minute step is considered completed when the loading equalisation between the two substations is reached to within a tolerance of $10^{-3}$, as stated in Algorithm 2. This relaxed tolerance is used because of the discretisation of the action space. Although this tolerance is larger than that typically used in an OPF, the overall objective is still achieved (i.e., a more efficient operation of the network) and for practical purposes it can be considered strict enough. Other termination states are also incorporated in the algorithm to ensure that certain constraints are respected by the agent. For instance, if the power rating limit of the converters is reached a large penalty is applied, the episode terminates and the agent has to start that episode again. Similar termination conditions are used to guard against a large increase in the power losses and time over-run (by imposing a maximum number of iterations permitted per episode of 2600). The information on whether a state is or is not a termination is encoded in the transition tuple. Finally, a whole day of operation (288 time-steps) is treated as an epoch and the agent is considered trained once an epoch is finished with all steps successfully converged. This self-learning algorithm presented in Algorithm 2.

**Algorithm 2** Self-learning algorithm for ANM using DQN

Initialize memory buffer to $N$ size
Initialize neural networks with random weights $\theta$, $\theta'$
Initialize LV network with load profiles for day $D$

**for** $minute = 1$ to $288$ **do**
  Solve power flow for LV network for $minute$
  Get initial state vector $s_t$
  **while** $|Tx_1S_{pu} - Tx_2S_{pu}| \geq 10^{-3}$ **do**
    Select action $a_t$ following $\epsilon$-greedy policy in (5.6)
    Assign $a_t$ to power converters
    Solve power flow for LV network
    Get $s_{t+1}$ and calculate reward $r_t$
    Evaluate termination statements
    Save transition $(s_t, a_t, s_{t+1}, r_t)$ to memory buffer
    Sample random batch of transitions from memory
    Calculate the target $y_t$ for each transition as in (5.10)
    Train neural network ($\theta$) with sample batch and $y_t$
    Every $C$ iterations copy weights $\theta$ to $\theta'$
    Decrease $\epsilon$ as in (5.7), until reaching $\epsilon_{min}$
  **end while**
**end for**

The overall model-free reinforcement learning algorithm is applied in ANM in two stages, an offline training stage where the self-learning agent is trained by interacting with the LV network model and an operational phase where the trained neural network is used to produce the desired set-points online. The deep neural network architecture and the continuous feedback received allow the agent to learn continuously and to update its internal representation of the characteristics of an LV network from the measurements, all without using an explicit model of the network. In contrast to the OPF where a new optimisation is required for every new operating condition, regardless of the similarity between one condition and another, the trained agent should be able to generalise to a certain extent to unseen states based on the learned transitions. The ability to generalise will be tested in the case study that follows by training the agent using load information from one specific day of the week and validating its performance on a different day.

### 5.2.3 Case Study - Lotherton and Canefields LV Substations

The performance of the proposed control algorithm is tested on a model of the two LV substations, Lotherton and Canefields, using historical data from the OpenLV project [55]. This case study was described in section 3.4.2 and again two converters with ratings of 90 kVA are used to connect the two substations as shown in Fig. 5.3. Only active power injections are considered. The control should provide a sequence of power injection set-points to the converters that minimise the power losses for the entire period analysed, in this case 288 time-steps, representing a day of operation.

The AC/DC OPF introduced in section 3.3 is used to compare and validate the performance of the self-learning approach. To make a direct comparison to the DQN algorithm, the objective function of the OPF will be to equalise the loading of the two LV substations. The results obtained by the OPF should act as an upper bound on the performance obtained by the self-learning agent, given its requirement for full-observability of the network parameters. It is worth noting that the optimal power injections from the converters depend on: the total impedance of the feeders, the load profile of the feeders, and the overall loading of the substations. However, both the OPF and the self-learning agent can only control the power at the two connected feeders, as seen in Fig. 5.3, and do not have any effect on the loads in any of the other feeders

A fully connected NN, with two hidden layers, each with 150 hidden units, was used to approximate $Q^{\pi}(s_t, a_t)$, as in (5.8). The neural network was trained using the Adam optimiser [173] with a learning rate of $10^{-3}$. On average, calculation of a single power flow in OpenDSS takes 50 milliseconds; hence, to analyse a minute of network operation with an average of 500 iterations, takes approximately 25 seconds. A training epoch (288 time-steps) for the agent took between 6 and 7 hours to complete using a desktop PC. This includes some episodes that required repeated attempts. Once trained, the control agent is validated by

loading the weights of the trained neural network and running the control on the same load profile. The time taken to complete the daily epoch in recall mode was much less than the training time. Loading equalisation was achieved in a single attempt for all episodes. To evaluate the model performance, the results of the self-learning algorithm were compared to those obtained by the OPF. Its performance was evaluated in terms of loading equalisation, total power losses and error with respect to the power injections calculated by the OPF. A summary of the parameters used for training is provided in Table 5.3.

Table 5.3: Summary of the parameters used for training.

| **Parameter** | **Value** | **Parameter** | **Value** |
|---|---|---|---|
| Hidden nodes | 150 | Activation | relu |
| Learning rate | 0.01 | Learning decay | 0.001 |
| $\tau$ | 0.9971 | $\epsilon_{min}$ | 0.01 |
| $\gamma$ | 0.95 | Batch size | 32 |
| Memory size | 100k | $C$ parameter | 30 |

Fig. 5.4 shows an example of training and covers the evolution of the converter's power injections and loading of the two substation transformers over 550 iterations during training for one point in time (time-step 200 of 288 steps and which corresponds to 16:40 hours). The transformers start with loadings of 53.13% for transformer 1 and 34.15% for transformer 2.. The algorithm then updates the converter power injections by accumulating the discrete actions from the possible set of actions, until loading equalisation is achieved to within a tolerance of $10^{-3}$. This takes 539 iterations. This example is typical of what was observed for all 288 time-steps.

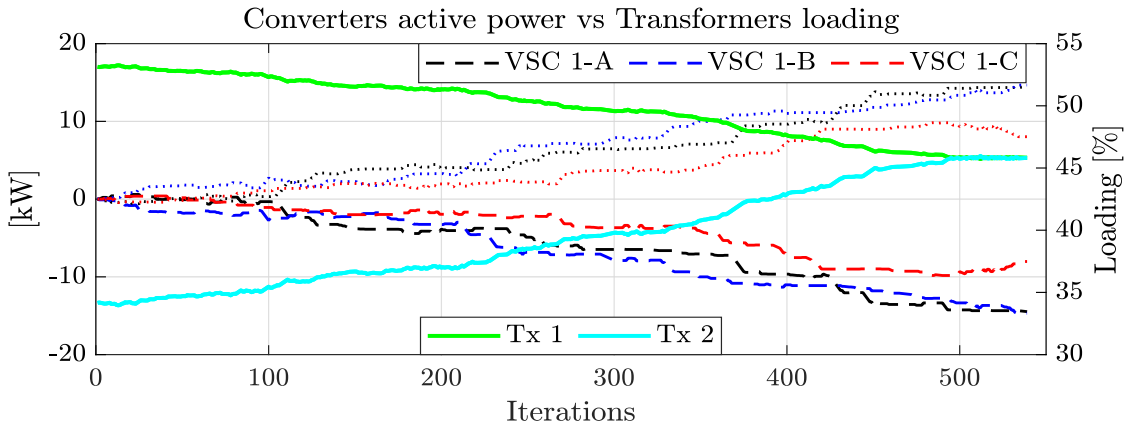



Figure 5.4: Evolution of the power injections of each phase of the two power converters and the loading of the two substation transformers which are to be equalised. These are the iterations at time-step 200. The dotted and dashed lines represent the converters' power injections, the solid lines indicate the transformers loading.

The cumulative reward for time-step 200 is shown in Fig. 5.5 and reveals a steady improvement as the algorithm advances. The power losses of the transformers are also seen to reduce during this process but it is worth noting that the final value of losses is not necessarily the minimum. This is because the self-learning

agent balances the goal of achieving loading equalisation and lowering the power losses, with the former taking priority. The evolution of the Q-values, depicted in Fig. 5.6, shows how they converge at an average value of 60, the maximum cumulative reward the agent can receive if it achieves the loading equalisation goal without increasing power losses.

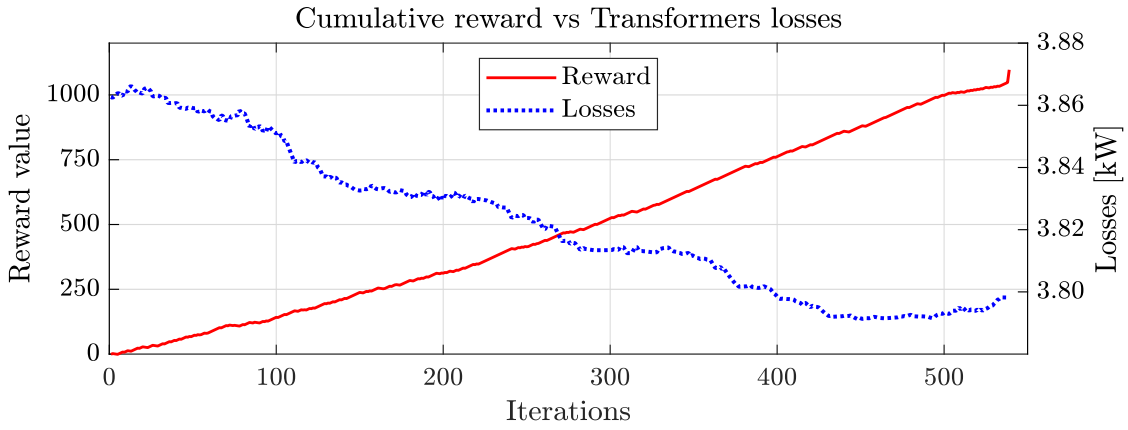



Figure 5.5: Evolution of the cumulative reward and the transformers' power losses for time-step 200.

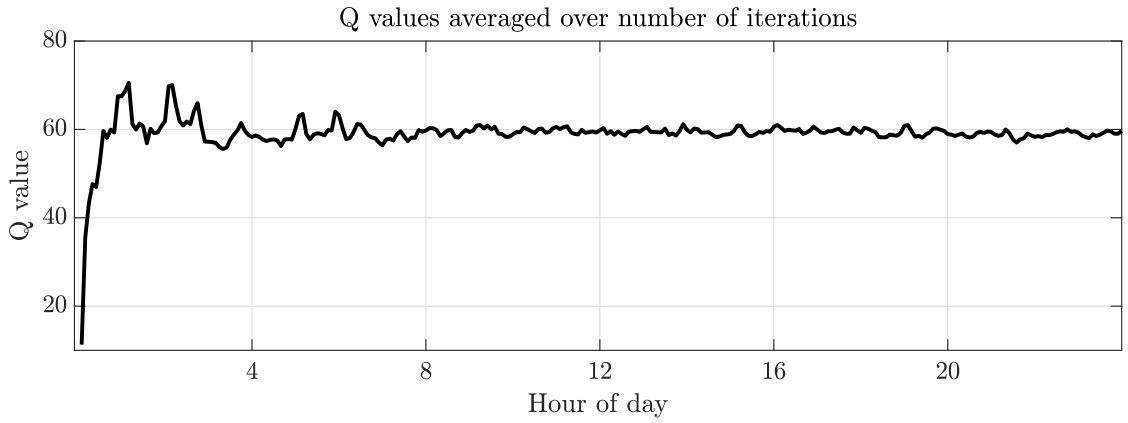



Figure 5.6: Q values of actions averaged over total iterations of each of the 288 time-steps.

A comparison of power injection set-points across all 288 time-steps of the OPF and the DQN is presented in Fig. 5.7. Notwithstanding a few discrepancies, the power injections calculated by the DQN reproduced successfully those of the OPF across the 24-hour period. More importantly, although there are multiple ways in which the load equalisation could have been achieved and several local minima might exist, the DQN algorithm appears to be in agreement with the OPF, and the objective is met in a similar manner.

To explore this further, a histogram of error between the set-points produced by the DQN and those of the OPF is displayed in Fig. 5.8. It can be observed that the error follows a normal distribution with a mean close to zero. Moreover, the majority of the differences are less than $\pm 4$ kW, with only a few occurrences greater than $\pm 6$ kW. These represents an error of 4% and 6% of the rated power of the converters. Both the root mean squared error (RMSE) and the mean absolute error (MAE) with respect to the OPF are calculated and are included in Fig. 5.8.

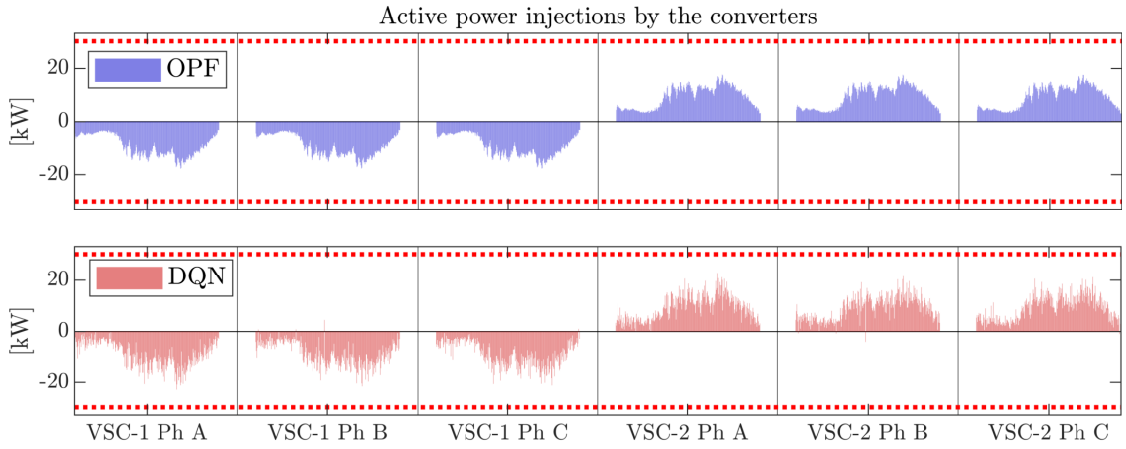



Figure 5.7: Comparison of active power set-points from the OPF and the DQN for each converter phase for Sunday (set out sequentially for the 24hr period).

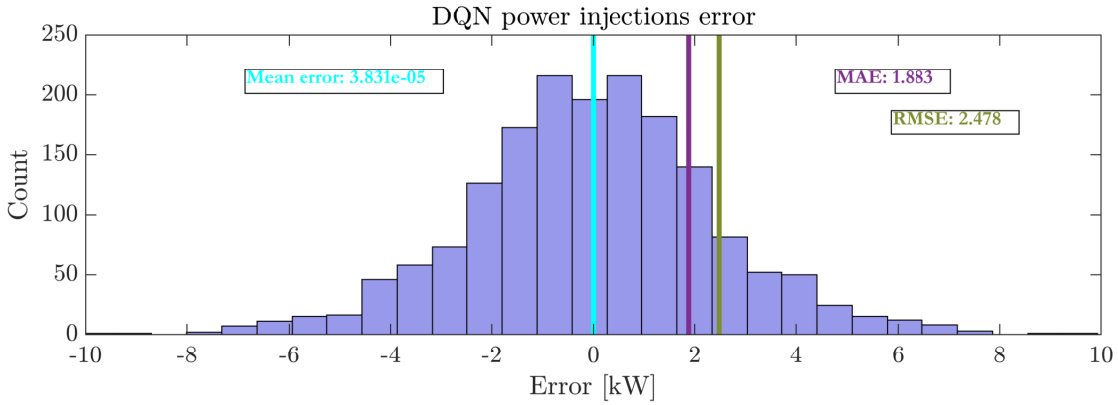



Figure 5.8: Histogram of the error between the injections calculated by the DQN and those by the OPF for the base case. The mean of the errors, the RMSE, and the MAE are included as annotations.

Fig. 5.9 shows that in the base case transformer 1 (Tx1) is more heavily loaded than transformer 2 (Tx2) throughout the day. The DQN control agent is seen to equalise the loading just as effectively as an OPF.

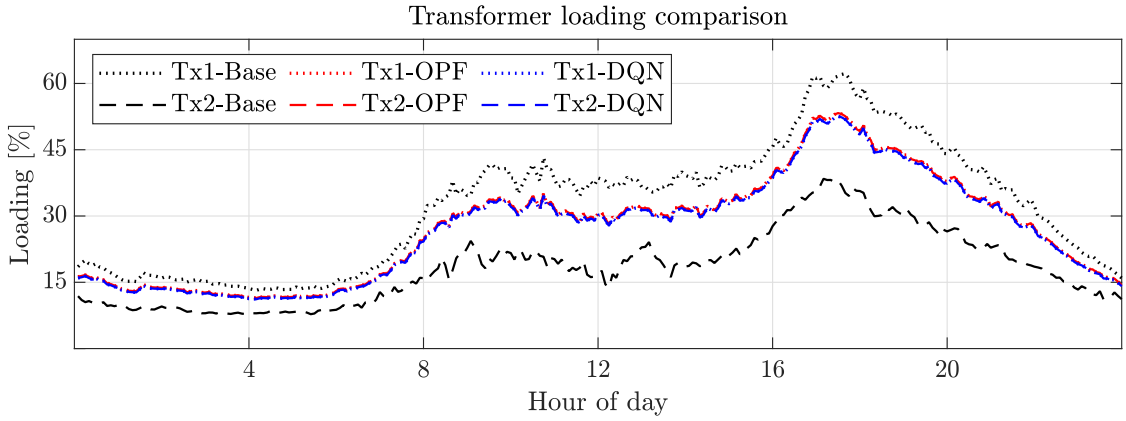



Figure 5.9: Comparison of transformer loading for the base case, equalisation using an OPF and the actions of the DQN control agent.

A comparison of the total power losses in the network is presented in Fig. 5.10. It can be seen that both the DQN and the OPF result in reduced power losses in the network compared to the base case where the converters were not used, particularly in the hours of peak load (17-19hrs), where the loading of the transformers is higher. Moreover, as observed in the Fig. 5.10 inset, the power losses results obtained by the DQN time-step closely follow those of the OPF. Although only a small decrease in power losses is observed, the main goal of keeping the loading of the transformers equalised is achieved.

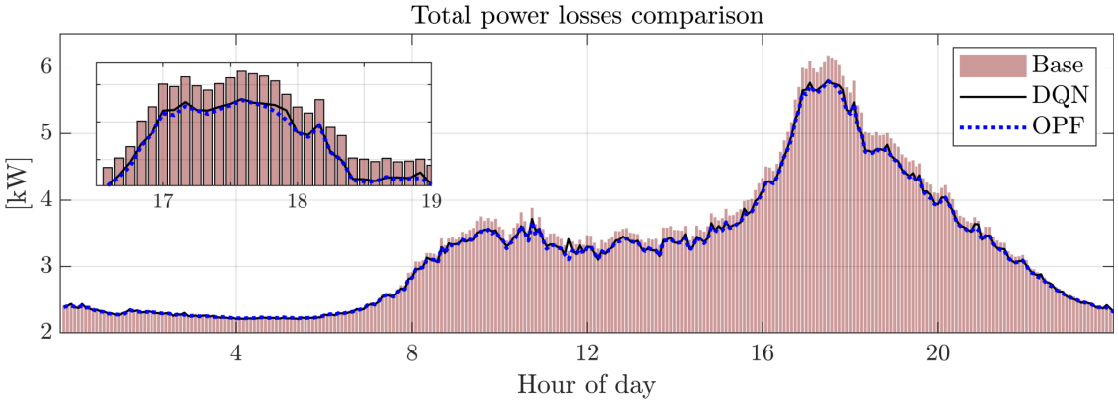



Figure 5.10: Total power losses comparison between base case, OPF, and DQN. The bars represent the power losses at each time-step for the base case; the solid and dotted lines represent the power losses for the DQN and OPF, respectively.

Finally, the results for the entire day are summarised in Table 5.4. Despite the OPF and the DQN agent using a different maximum power injections, the overall benefits obtained by the DQN in terms of power losses and loading equalisation are similar to those of the OPF. The total power losses are reduced and the utilisation of the transformers is decreased by equalising the loading across the LV substations. An important consideration is the fraction of transformer rating (apparent power) that remains available for further use

(hosting LCTs for instance) at the peak hour. The DQN agent achieves a larger combined remaining capacity of 387.1 kVA than the 382.6 kVA present in the base case.

Table 5.4: Results comparison for entire day - Training.

| Case | Max power transfer [kW] | Losses network [kWh] | Tx 1 + Tx 2 Losses [kWh] | Tx 1 Loading max [%] | Tx 2 Loading max [%] |
|---|---|---|---|---|---|
| Base | - | 80.08 | 67.16 | 62.3 | 38.4 |
| OPF | 17.77 | 77.18 | 66.18 | 53.3 | 53.3 |
| DQN | 22.77 | 77.46 | 66.12 | 52.5 | 52.5 |

The ability of the trained agent to generalise is validated by evaluating its performance on a load profile of a different day; for example, Monday versus Sunday. This is done to ensure the agent does not memorise the load profiles but learns a robust control policy that is able to generalise to small variations in the load profiles.

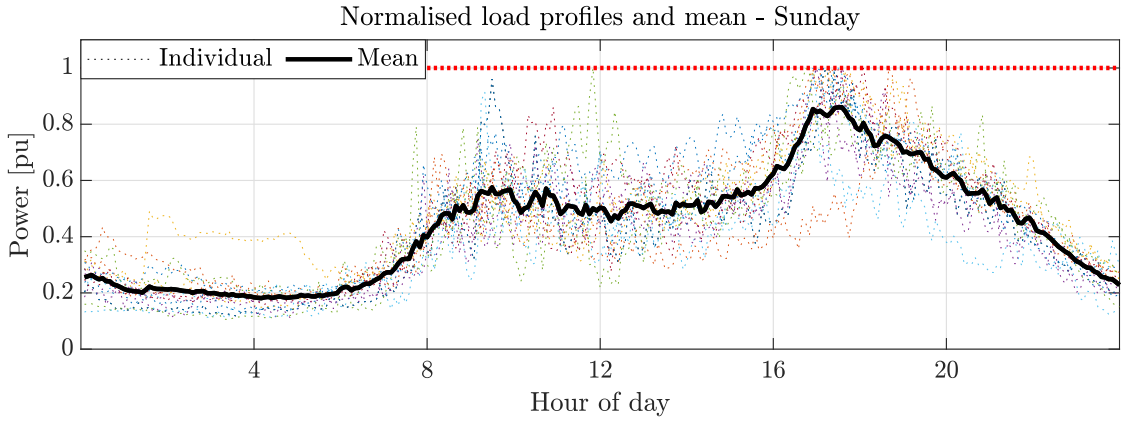



Figure 5.11: Load profiles at the head of the feeders and their means for Sunday.

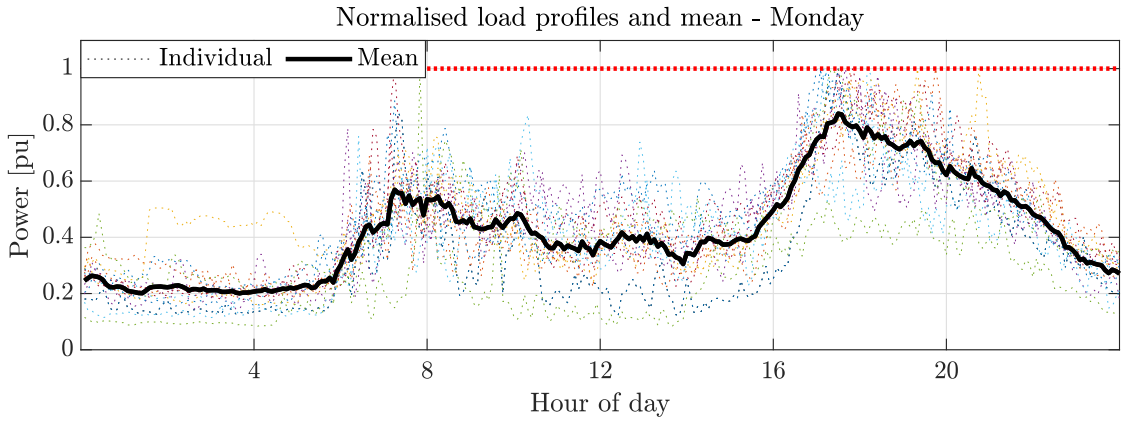



Figure 5.12: Load profiles at the head of the feeders and their means for Monday.

Fig. 5.11 and Fig. 5.11 shows a comparison of the training (Sunday) and validation (Monday) load profiles, and their averages, used for both substations. Considerable differences are observed in the load profiles. For

instance, the morning peak starts earlier and has a larger value for Monday, whereas the loading at noon is higher for Sunday.

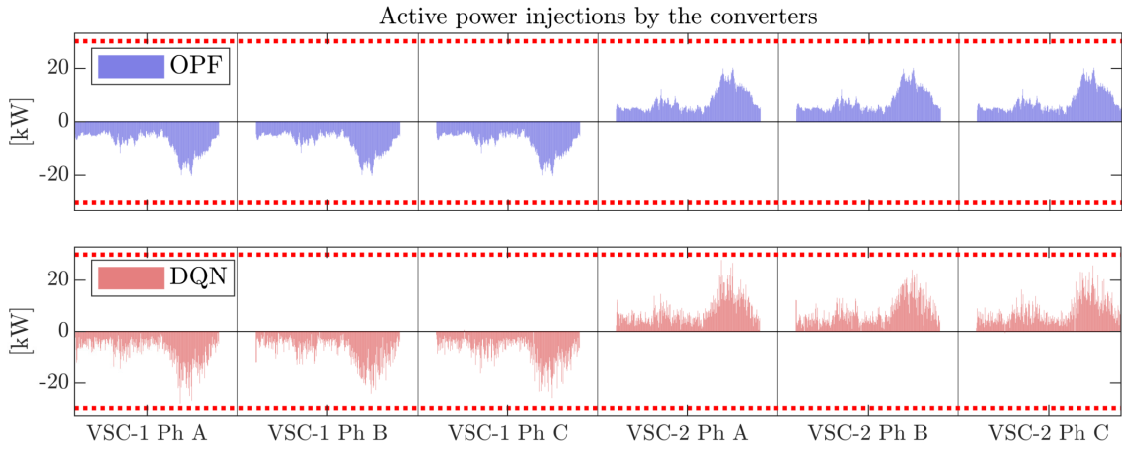



Figure 5.13: Comparison of active power set-points from the OPF and the DQN for each converter phase for Monday (set out sequentially for the 24hr period).

The power injections obtained by the trained DQN agent and the OPF, for the Monday profile, are shown in Fig. 5.13. Similar to the results of the training, the power injections obtained by the DQN are very close to the OPF, with only a few deviations.

The respective error is displayed in Fig. 5.14. The distribution is close to normal centred around zero and both the MAE and RMSE have acceptable values of 2.052 kW and 2.682 kW, respectively. Fig. 5.15 confirms that the self-learning agent, trained on Sunday data, can still successfully equalise the transformers' loading for all 288 time-steps of the Monday data.

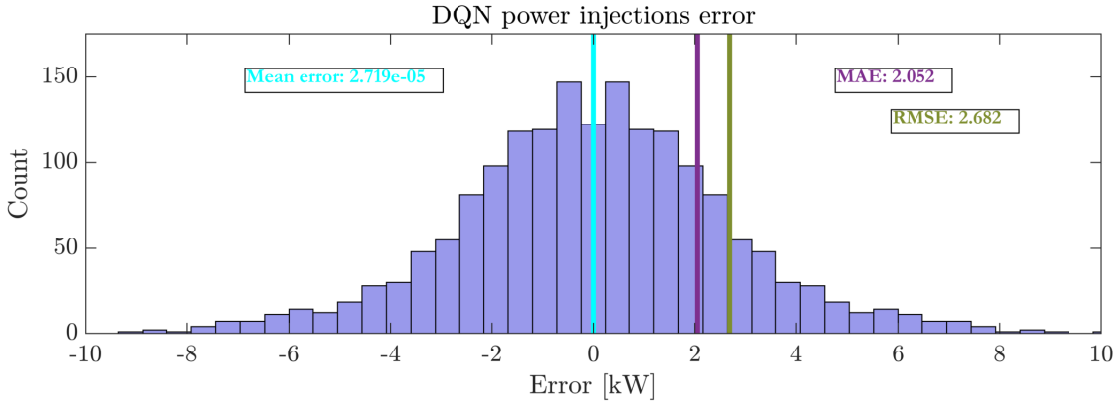



Figure 5.14: Histogram of the error between the injections calculated by the DQN and those by the OPF for the Monday load profile. The mean of the errors, the RMSE and the MAE are included as annotations.

Table 5.5 summarises the validation results for the generalisation test. In this case, the overall capacity available at the peak hour for both transformers is increased to 396.1 kVA, versus 394.0 kVA for the base case. It is worth highlighting that the benefits of the converters together with a self-learning control would

increase as more LCTs are connected. Also, a larger amount of capacity could be released if the substations had a complementary load profile.

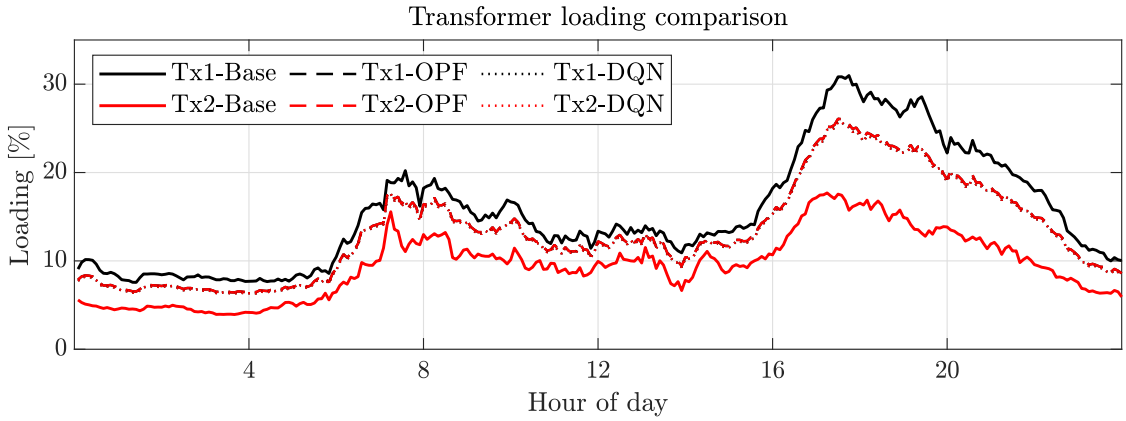



Figure 5.15: Transformers' loading comparison between the base case, DQN, and the OPF for the Monday load profile.

Table 5.5: Results comparison for entire day - Validation.

| Case | Max power transfer [kW] | Losses network [kWh] | Tx 1 + Tx 2 Losses [kWh] | Tx 1 Loading max [%] | Tx 2 Loading max [%] |
|---|---|---|---|---|---|
| Base | - | 77.95 | 65.82 | 61.9 | 35.4 |
| OPF | 20.26 | 75.59 | 65.10 | 52.2 | 52.2 |
| DQN | 27.87 | 75.90 | 65.05 | 51.4 | 51.4 |

In the case of DQN, reusing the converters' set-points from previous episodes (i.e. from the preceding minutes) did not lead to a good performance, as the set-points obtained in the preceding episodes might result in a violation of the termination statement, terminating the episode. This would make the training might become too inconsistent, leaving the agent with no chance to explore for some iterations and improve its learning. Treating each minute independently, as a separate training episode, allows the agent to explore with the epsilon-greedy policy, in order to store batches of experience with random actions. To allow this the epsilon value was reset to $\epsilon = 1$ (exploration) at the beginning of each minute. This way the agent would get the chance to try a set of completely different random actions.

## 5.3 Model-based Reinforcement Learning for ANM

A different form of RL is possible when a full model of the network can be built and the equations describing the system to be controlled are available; in the ANM case these would be the power flow equations representing a given network and the injections from the converters. With the equations available, the learning process can benefit from evaluating the derivatives of these equations with respect to any variable of

interest [191]. These derivatives, or gradients, can indicate the sensitivity of the cost function to the inputs and control actions. Thus, the control actions can be directly optimised through back-propagation and gradient-based optimisation, as proposed in [198, 199]. This optimisation process is analogous to the neural network training procedure described in section 4.2 and is widely used for training deep neural networks for tasks ranging from computer vision to natural language processing [200]. However, in this case, instead of minimising the error between the targets and the neural network output, an optimal schedule of actions is obtained by iteratively updating the actions based on the gradient information. For the optimisation to make use of the gradients, the model representing the system must be differentiable [172, 201] and having a sparse structure enables a more efficient calculation of the derivatives. If network components that have discrete states, such as tap-changers and switched capacitor banks, are ignored the power flow equations are differentiable and can be readily used. Moreover, the admittance matrix, representing the connections between nodes of an LV network, is inherently sparse [83, 110]. Using this type of model-based RL framework becomes advantageous when the computational cost of learning the model is much lower compared to finding an optimal policy or series of set-points [201]. This is usually true in the case of models of typical LV networks, for which calculating its state at any given point via a power flow time-step is computationally cheap compared to performing optimisation via an OPF. This is indicated by the computation times found in the case studies in section 3.4. For this reason, a significant speed improvement is anticipated from this model-based framework compared to the OPF.

Back-propagation, as introduced in [172, 199, 202], uses the chain rule for *ordered derivatives* to calculate the exact derivatives of any nonlinear, sparse system in a single pass. The system can be composed of several subsystems, represented by an ordered set of equations. Unlike conventional partial derivatives, *ordered derivatives* represent both the direct and indirect impact of an action on a system, that is, they take into account the effects of an action on all of the variables comprising a system [172]. As described in section 4.1, back-propagation is a general purpose algorithm that can be used to calculate the gradient of input-output relationships by propagating operations through the system. Back-propagation, as a tool for calculating gradients, can be seen as a special case of the reverse accumulation mode of automatic differentiation [166]. An extension known as back-propagation through time (BPTT) was proposed in [172, 198] as a control scheme for dynamic systems. One of the main differences of BPTT is that, long-term effects of the control variables are accounted for in the cost function being maximised [172]. Thus BPTT can be seen as a dynamic optimisation over time, where the output is not only a set of optimal points but an optimal trajectory. In contrast to the OPF, which optimises each time step separately, the method proposed in this section takes into account the long term effects and can be seen as a dynamic or multi-period optimisation. This optimisation framework can be extended further to consider the case where a neural network is trained to take these control actions. This type of control is part of the framework proposed by ADP [203, 204] and can

be implemented using a neural network as an action network. In this case, the system to be controlled is represented as a differentiable model and rather than optimising the actions directly, a neural network, an *actor*, is trained to take optimal actions based on the direction of the gradient of the output being optimised with respect to the neural network weights. The model of the system, the *critic*, serves as feedback to train the neural network to choose optimal actions. The term *critic* was first coined by Widrow in [205], to refer to a control component that evaluates the decisions made with respect to a perfect model.

### 5.3.1 Tools for Automatic Differentiation

The TensorFlow library enables the efficient construction of *computational graphs* composed of a set of nodes representing the variables and operations, and it facilitates the calculation of the derivatives of these graphs [206]. To do this, it uses the reverse mode accumulation of automatic differentiation (AD) [206, 207]. Unlike numeric and symbolic differentiation, AD does not provide a numerical approximation of the derivatives, nor does it provide explicit derivative expressions [207–209]. Instead, it combines features from both numeric and symbolic differentiation to deliver an exact but computationally efficient evaluation of the derivatives of an expression. The reverse mode of AD is able to calculate the derivatives of an expression, with respect to any of the variables, in two passes through the graph. This is because each partial derivative is stored as an intermediate step in the computational graph. In the forward pass, the value of each intermediate operation is calculated, in the second pass (backward direction) all the partial derivatives are calculated and propagated from the outputs to the inputs [206, 207]. Hence, for expressions with a large number of inputs and a small number of outputs the reverse mode exhibits a greater computational efficiency, compared to the forward mode [206, 207].

One of the main advantages of AD (as implemented in TensorFlow) is that, unlike symbolic differentiation, AD does not require closed-form expressions but can also work with expressions in the form of general algorithms. For instance, it can be used in combination with complex expressions involving *if-then-else* decisions and *while* loops (also referred to as control-flow operations) [206]. Thus, power flow algorithms, such as the fixed-point algorithm presented in section 2.2, can be readily implemented as a computational graph. This enables the efficient calculation of the gradient between output and input variables, ultimately enabling the optimisation of variables of interest through gradient descent algorithms. This will be demonstrated in a case study in the next section.

AD has been developed from the fact that any complex algebraic functions can be split into a combination of several elementary operations (e.g. additions, multiplications, trigonometric functions, etc.) [166, 192, 206]. These elementary operations can be organised into a computational graph, with multiple input/output variables, which can be evaluated and differentiated *programmatically*. The computational

graph, composed of nodes and edges, stores the dependencies between variables or inputs/outputs for each of the elementary operations comprising the final expression. Because it tracks all the intermediate operations (and its respective derivatives), it is able to calculate how a change in each variable affects the final expression [168,192,207]. This is useful for determining the sensitivity of an action on a network parameter, such as the voltage or the power losses. Nonetheless, a side effect of this is that the computational cost increases with the number of intermediate operations part of the final expression. This is illustrated in the case study of section 5.3.3, where a number of fixed-point iterations are needed for the power flow to converge to a solution at each time-step of the daily loading profile.

An example of a simple computational graph can be observed in Fig. 5.16. Where the inputs are the variables $Z$ and $Y$, in blue, and the constant 2 in red. The output of this graph, in yellow, is obtained by performing the elementary operations on the inputs. Fig. 5.17 illustrates how a computational graph is built by TensorFlow and visualised in TensorBoard.

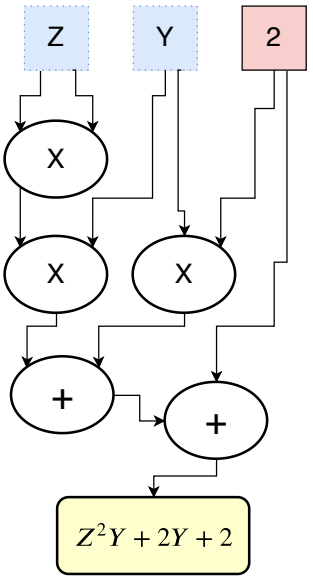



Figure 5.16: Example of a simple computational graph with two variables and one constant [210].

The construction of computational graphs facilitates the efficient computation of the derivatives of operations. Fig. 5.18 shows a computational graph, and its auxiliary variables, numerically evaluated in the forward direction (forward pass) at a specific point ($X = 3$, $Y = 1$ and $Z = 2$). The reverse mode (backward pass) of AD can be used to obtain the gradient with respect to the variables as follows:

The partial derivatives of the intermediate steps are calculated in (5.11), (5.12) and (5.13)

$$\frac{\partial d}{\partial c} = a, \ \frac{\partial c}{\partial b} = \cos(b), \ \frac{\partial b}{\partial Z} = Y \tag{5.11}$$

$$\frac{\partial d}{\partial a} = c, \ \frac{\partial a}{\partial Y} = X, \ \frac{\partial a}{\partial X} = Y \tag{5.12}$$

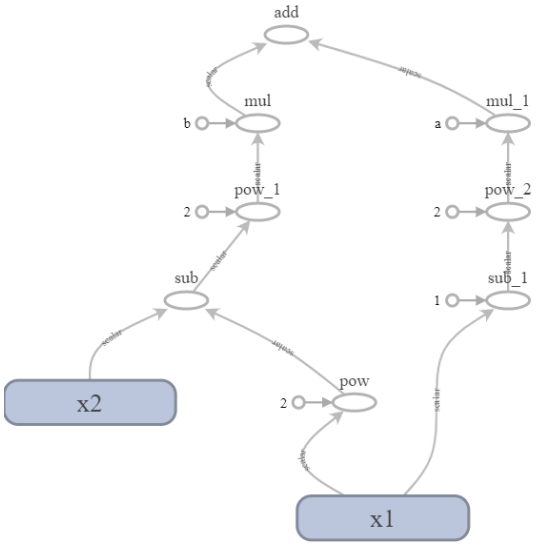



Figure 5.17: Example of a computational graph in TensorFlow.

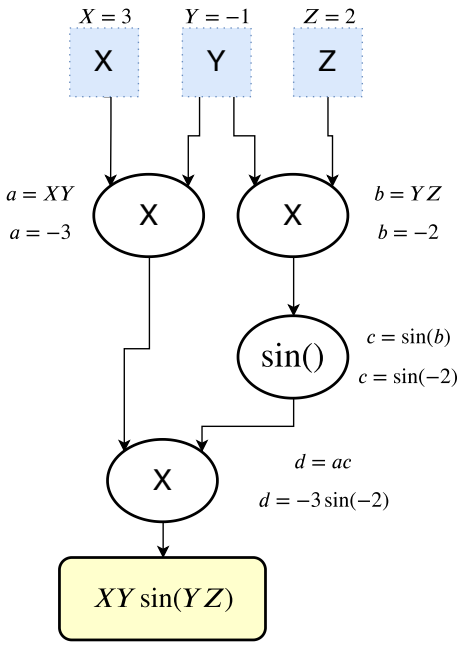



Figure 5.18: Computational graph with intermediate results to illustrate the reverse mode of AD [211].

$$\frac{\partial d}{\partial X} = \frac{\partial d}{\partial a}\frac{\partial a}{\partial X} = -\sin(-2) \tag{5.13}$$

The gradient of the inputs $Y$ and $Z$, with respect to the final output, can then be calculated as in (5.14) and (5.15). For a detailed description of the two modes of AD and their implementation in TensorFlow, the reader is referred to [209], [206] and [207].

$$\frac{\partial d}{\partial Y} = \frac{\partial d}{\partial a}\frac{\partial a}{\partial Y} + \frac{\partial d}{\partial b}\frac{\partial b}{\partial Y} = \frac{\partial d}{\partial a}\frac{\partial a}{\partial Y} + \frac{\partial d}{\partial c}\frac{\partial c}{\partial b}\frac{\partial b}{\partial Y} = 3\sin(-2) - 6\cos(-2) \tag{5.14}$$

$$\frac{\partial d}{\partial Z} = \frac{\partial d}{\partial b}\frac{\partial b}{\partial Z} = 3\cos(-2) \tag{5.15}$$

### 5.3.2 Adaptive Neural Networks for ANM

The model-based algorithm proposed in this section follows the power flow formulation introduced in section 2.2 combined with a fixed-point algorithm to calculate the vector of node voltages and other network variables. In particular, it uses the Algorithm 1 (repeated here as Algorithm 3 for ease of reference) in the form of a computational graph to calculate the power flows of the network in question. As it can be observed from 3, the fixed-point requires a number of iterations for the voltage to converge within the error threshold.

Unlike the OPF used in section 3.3, this optimisation approach does not require the analytical evaluation, or numerical approximation, of the Jacobian (first-order) and Hessian (second-order) derivatives, which are required in an OPF, such as the one using the IPM and where a function such as MATLAB's *fmincon* might be used. Instead, it makes use of TensorFlow's native AD engine to calculate the exact derivatives at any point. Thus it can calculate the gradient of the network power losses with respect to the power injections of the converters. Using the gradient, the algorithm can iteratively adjust the power injections to minimise the power losses through gradient descent. As noted before, the long term effects are considered by this method, thus the impact of the converters on the global power losses (i.e. the losses throughout the day) is considered.

As described in 5.3, on a first instance, the actions can be optimised directly with gradient descent, without a neural network, making this model-based RL algorithm a computationally efficient alternative to the OPF. The first step is to build a computational graph, which comprises the required iterations for the fixed-point algorithm to converge to a feasible solution within the established error tolerance ($V_{error}$), as detailed in (2.28). The next step calculates the gradient of the total power losses with respect to the power injections of the converters. Finally, this gradient is used to optimise the power injections through gradient descent, or an extension of it (e.g. gradient descent accelerated by momentum [173]), acting directly on the variables

**Algorithm 3** Fixed-point algorithm for power flow

Get primitive admittance matrices ($Y_{prim}$) of each element in the network

Get nodal admittance matrix of the system ($Y_{system}$)

Get voltage of the equivalent source at the substation assuming nominal voltage $V_{base}$ and balance

Get injection current from the substation $I_{source}$

Start iterations $k = 0$

Initial guess at the voltage vector $V_{nodes}^{(k-1)}$ without the loads compensation currents as in (2.26)

**while** $V_{error} \geq 10^{-5}$ **do**

    Update the iterator counter $k = k + 1$

    Calculate linear part of current for the loads as in (2.22)

    Calculate injection current for the loads as in (2.23)

    Calculate compensation current for the loads as in (2.25)

    Update the voltage vector $V_{nodes}^{(k)}$ including the compensation current of the loads as in (2.26)

    Calculate the $V_{error}$ as in (2.28)

**end while**

Calculate the current in all lines with the final $V_{nodes}$ and the respective $Y_{prim}$ as in (2.19)

Calculate the power flow in all lines with the final $V_{nodes}$ and the respective currents as in (2.20)

of the computational graph. The optimisation process is run for a number of training iterations until it converges to a local minimum. An overview of this process is shown in Fig. 5.19.

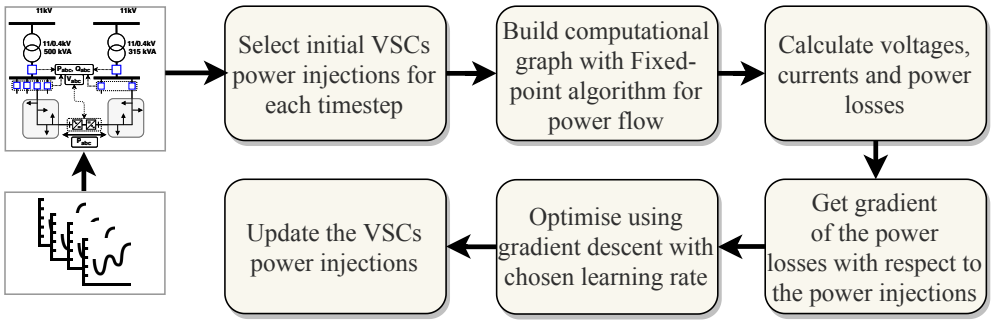



Figure 5.19: Overview of the model-based RL optimisation environment. The network and the load profiles are modelled using the power flow equations, which solved via the fixed-point iteration method. The actions are optimised directly by gradient descent.

The model-based algorithm can then be extended by incorporating a neural network to choose the actions (i.e., the action network or actor) into the algorithm. Because of this, the back-propagation and optimisation are now realised with respect to the weights of the neural network, instead of the actions directly, as described in section 5.3. An overview of this stage is shown in Fig. 5.20.

By training a neural network, the algorithm is able to deal with partial observability, as only some of the

network measurements are made available as inputs to the neural network. Even though the optimal power injections are obtained using an explicit model of the LV network, the neural network that produces these decisions (the actor) is only allowed to *observe* a subset of the measurements, those that are normally available to DNOs through their instrumentation and data-acquisition infrastructure.

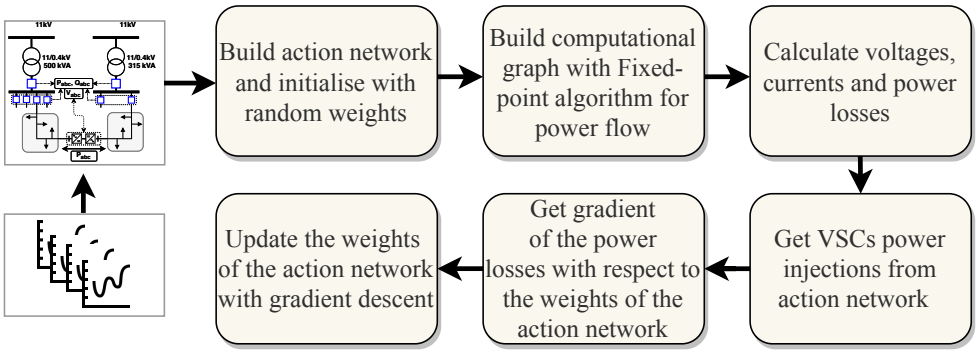



Figure 5.20: Overview of the model-based optimisation environment, similar to Fig. 5.19 but including the action network, which produces the actions.

The *observable* state vector can incorporate variables such as: the three-phase current from the transformer ($I_{Tx}$) and the voltage at the terminals of the converters ($V_{VSC,1}, V_{VSC,2}$). Additionally, since the power losses at each of the individual lines along a feeder cannot be measured in practice, the transformer losses, which depend on the total current delivered by the substation, are also used as input to the neural network.

### 5.3.3 Case Study - IEEE 123-bus Test Feeder - no Action Network

In order to compare the performance of the proposed algorithm, the IEEE 123-bus test feeder is used here as a case study. Solar PV generation and EVs are included, similar as in section 3.4.1. Also, the same objective function is chosen, to minimise the total active power losses in the network, to ensure a direct comparison can be made. Gradient descent with a learning rate of 0.65 was used to optimise the neural network. A total of 20 training iterations were performed in this case study. An informal search with different learning rates was made but 0.65 resulted in the best balance between total number of iterations and local minimum reached. Hence, the results are reported for this learning rate only.

The evolution of the total daily losses at each training iteration is shown in Fig. 5.21. It can be seen that the first steps of the gradient descent reduce the losses significantly from the baseline 1943.89 kWh, and the progress then slows down until it converges to a local minimum of 1359.96 kWh. This is very close (less than 1% difference) to the minimum reached by the OPF 1355.80 kWh, in section 3.4.1. The breakdown of the active and reactive power losses by component type is shown in Fig. 5.22 and can be be compared with the results from the baseline and the OPF in Fig. 3.11 and Fig. 3.22, respectively.

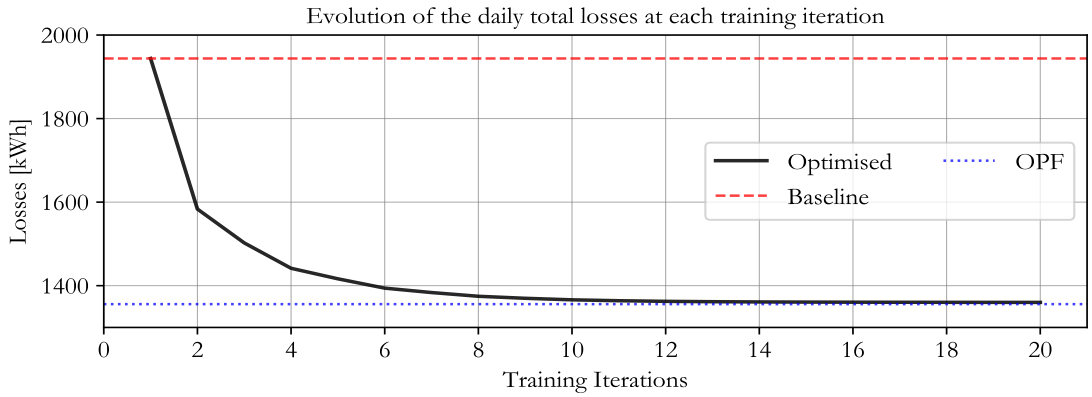



Figure 5.21: Total daily energy losses at every training iteration. The baseline is shown by the red dashed line whereas the minimum reached by the OPF is shown by the blue dotted line.

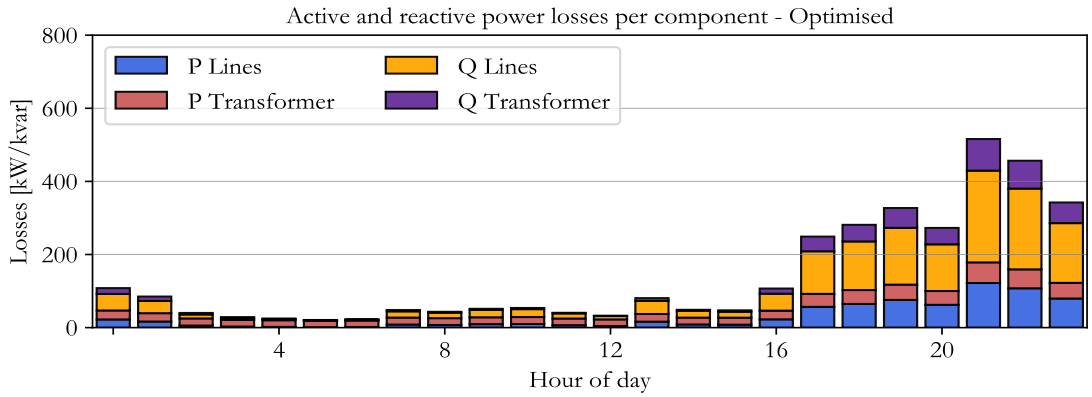



Figure 5.22: Active and reactive power losses by component type.

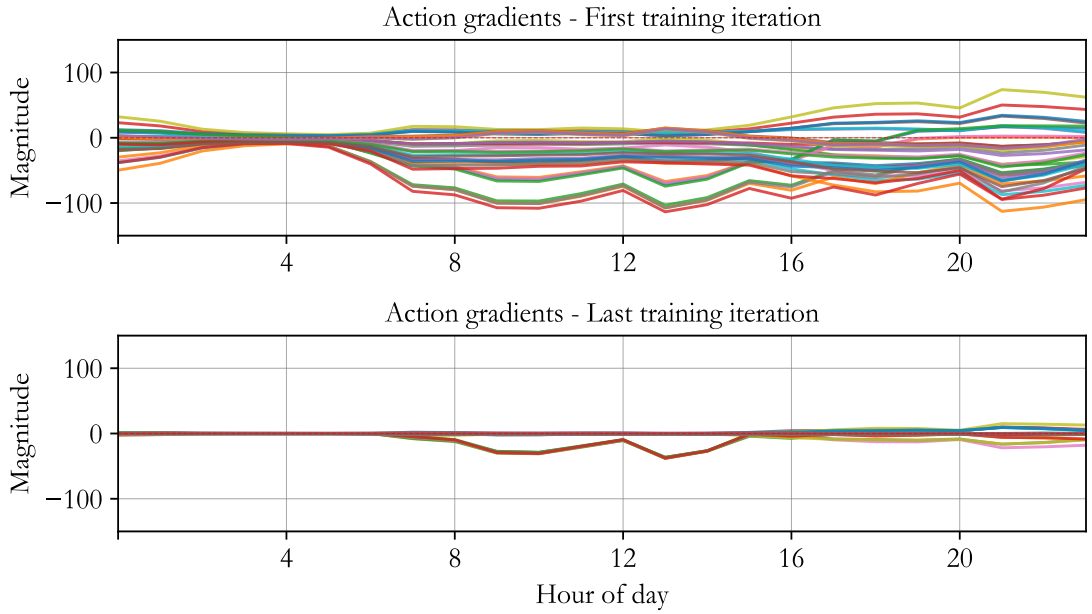



Figure 5.23: Value of the gradients during the first and last training iteration. Each of the colours represent the gradient of the total power losses with respect to one of the actions (i.e., the converters power injections).

The value of the gradients of the cost function with respect to each of the actions, during the first and last training iteration, are shown in Fig. 5.23. It can be seen that overall the value of the gradients is decreased significantly, with the majority of them reaching a value close to zero. This is another indication that the gradient descent algorithm converged satisfactorily to a local minimum.

The active power injected by the converters is shown in Fig. 5.24. Some differences are observed between the injections obtained here an those in Fig. 3.18. For instance, the power injections of VSC 2 phase C and VSC 4 phase A are larger for the OPF case than the model-based RL case. In the case of VSC 3 A and VSC 4 phase B the injections have a different direction.

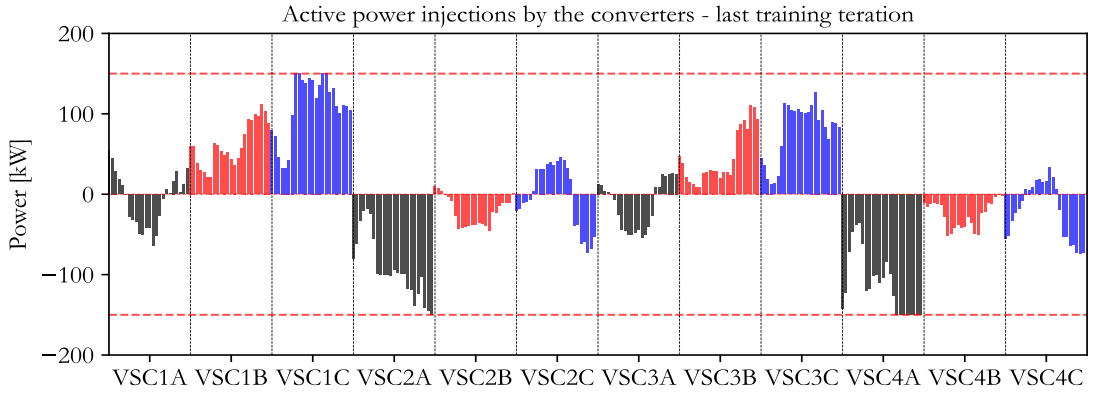



Figure 5.24: Active power injected or absorbed by converters set out sequentially for the 24hr period. The different colour of the bars indicate the phases. The red dashed line indicates the maximum active power that the converters can inject.

Fig. 5.25 shows the case of the reactive power, where it can be seen that the differences are smaller compared to Fig. 3.19.

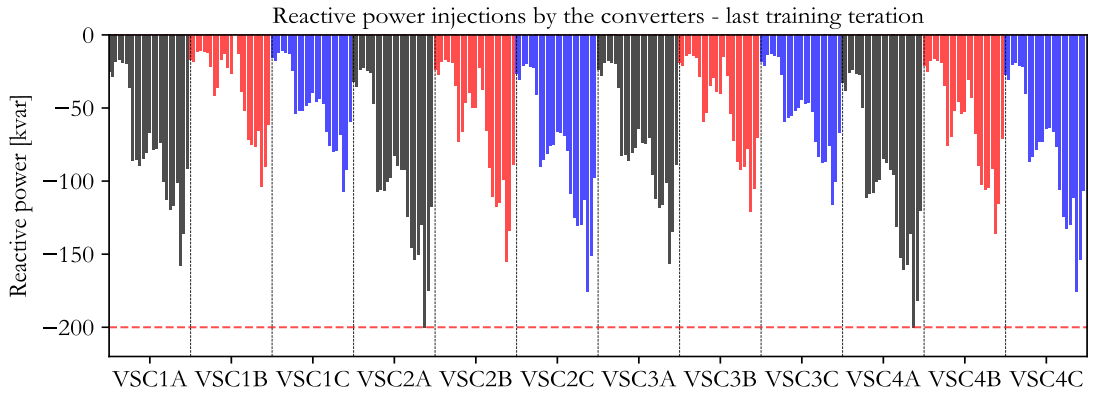



Figure 5.25: Reactive power injected or absorbed by converters set out sequentially for the 24hr period. The different colour of the bars indicate the phases. The red dashed line indicates the maximum active power that the converters can inject.

Fig. 5.26 shows a comparison of the current supplied by the substation for the baseline and at the end of the optimisation. Similar to the OPF result in section 3.4.1, the high phase imbalance, leading to a large

neutral current, is reduced greatly, especially during the hours of peak load and peak solar PV generation.

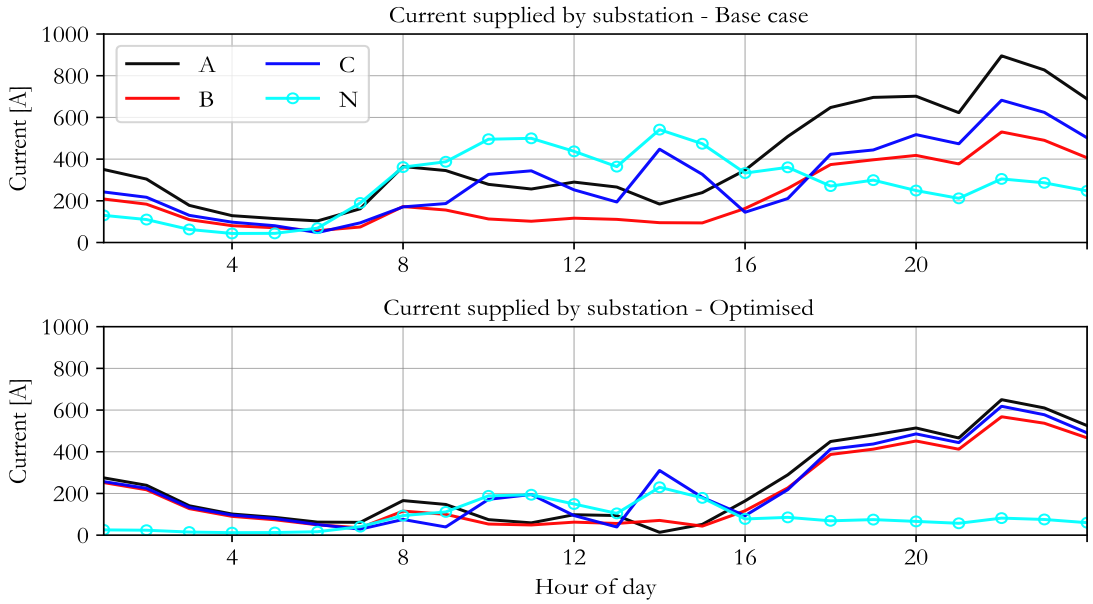



Figure 5.26: Comparison of the current supplied by the substation, for the baseline and for the optimised case.

A comparison between the power set-points resulting from the model-based optimisation and those of the OPF is shown in Fig. 5.27. It can be seen that despite some differences, overall there is agreement between the two solutions.

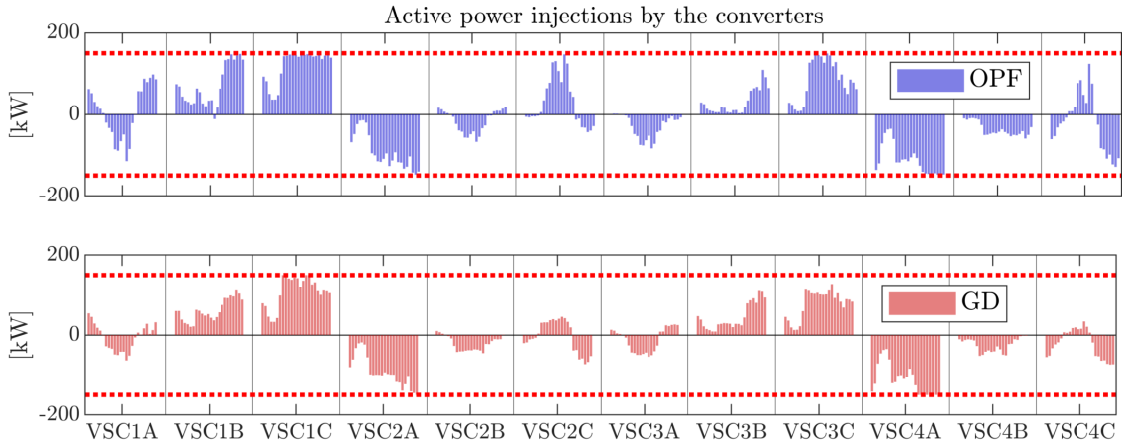



Figure 5.27: Comparison of active power set-points from the OPF and the Gradient Descent optimisation for each converter phase (set out sequentially for the 24hr period).

The difference or error between the set-points produced by the Gradient Decent optimisation and those of the OPF is displayed in Fig. 5.28, together with the RMSE and the MAE. Although the distribution is not strictly normal, with a slight skew towards the negative values, the mean of 2.815 kW is very close to zero. The mean difference represents less than 2% of the rated power of the converters, with some outliers (less than 10 occurrences) with differences of up to $\pm 50$ kW, which represents an error of 20%. It is worth

emphasising that similar to the model-free case, a complete agreement is not expected as the two optimisation frameworks can reach different local optimum. Nonetheless, although the optimal set-points obtained by Gradient Descent can be different depending on factors such as the learning rate, it was demonstrated in Fig. 5.21 that the local minimum reached is within 1% of that obtained by the OPF.

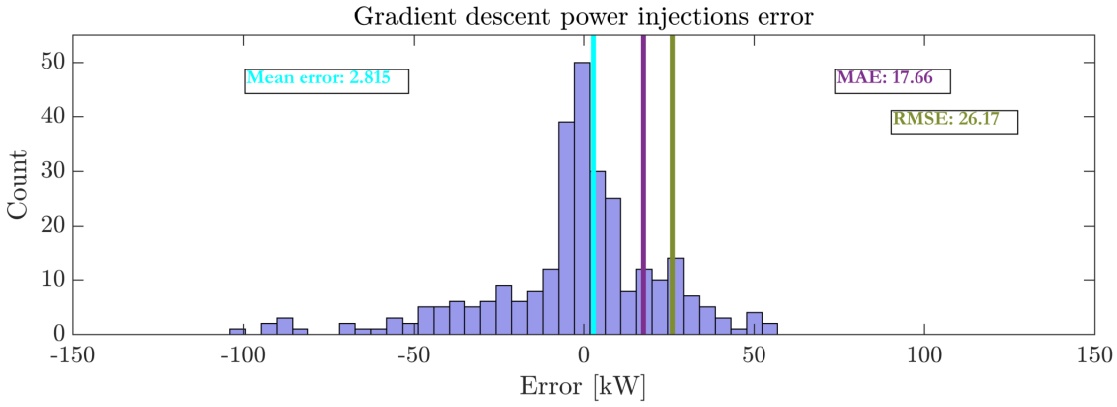



Figure 5.28: Histogram of the error between the injections calculated by the DQN and those by the OPF for the Monday load profile. The mean of the errors, the RMSE and the MAE are included as annotations.

Regarding the computational runtime, for this particular case study, building the computational graph takes on average 130 seconds, setting up the gradient operation takes around 450 seconds. The 20 training iterations take on average 20 minutes. The first iteration of the training procedure takes longer (700 seconds) than the subsequent (30 seconds each on average) because of TensorFlow's just-in-time compilation approach. Fig. 5.29 summarises the runtime taken by each operation. Overall the optimisation of the 24 time-steps takes approximately 35 minutes, which represents less than a third compared to the 146 minutes OPF runtime in section 3.4.1. Additionally, it is worth noting that due to the computational graph structure, this approach could be implemented in a graphics processing unit (GPU) or a Tensor Processing Unit (TPU), speeding up the computational time even further.

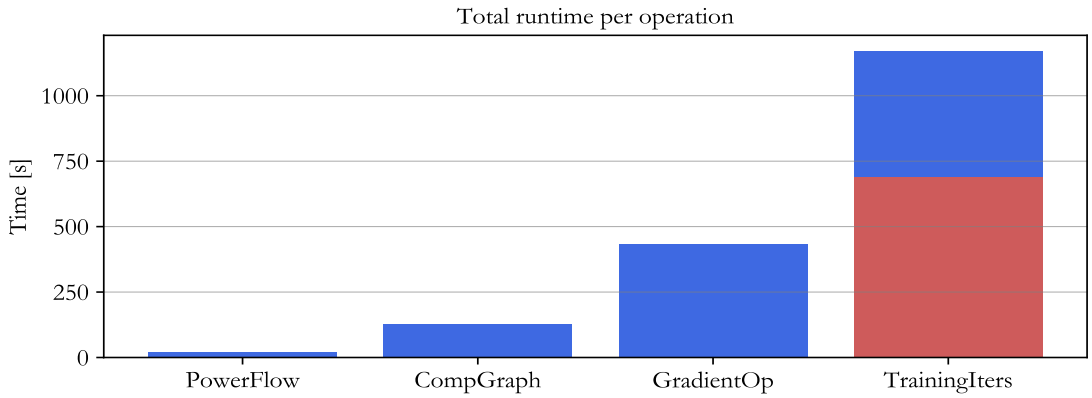



Figure 5.29: Runtime for each operation. For the training iterations, the red bar indicates the time taken by the first training iteration, the blue bar indicates the time taken by the subsequent 19 iterations.

Also, it is worth highlighting that both the construction of the computational graph and setting up the gradient operations are only performed once and remain *static*, TensorFlow does not execute the operations comprising the graph unless it is directed to. This means that unless the admittance matrix of network changes due to physical changes or reconfiguration this computational overhead (CompGraph and GradientOp in Fig. 5.29) only has to be done once for each LV network. The variables to be optimised (converter power injections) and the network loading/generation profiles can be modified to initiate a new training operation (TrainingIters in Fig. 5.29) and obtain a new series of optimal set-points from the *static* computational graph.

### 5.3.4 Case Study - IEEE 123-bus Test Feeder - with Action Network

The results of the action-network approach, described in Fig. 5.20, are shown in this section. A neural network with 18 inputs, 36 outputs and 2 hidden layers, each with 72 nodes, was used as the action network, as illustrated in Fig. 5.30.

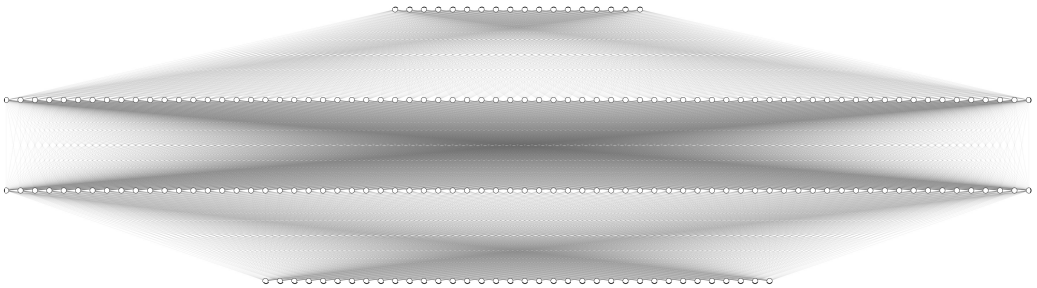


Figure 5.30: Illustration of the neural network architecture, comprising 18 inputs (at the top), 36 outputs (bottom) and 2 hidden layers [170]. This neural network has more than 9,000 trainable weights.

Due to the more complex optimisation involving the gradient of the losses with respect to the weights of the neural network, 25 training iterations were required to converge to a local minimum. In this case, the Adam optimiser from [173] was used, with a learning rate of 0.055.

The evolution of the total daily losses at each training iteration is shown in Fig. 5.31, where it can be observed that in this case the optimisation does not follow a smooth path as in Fig. 5.21. However, the algorithm converges to a local minimum of 1402.1 kWh, with a difference of 3% compared to the minimum reached by the OPF 1355.80 kWh. This compromise is acceptable given that only limited information is made available as input to the neural network. The active and reactive power injections by the converters are shown in Fig. 5.32 and Fig. 5.33. Some differences are observed compared to Fig. 5.24 and Fig. 5.25 but overall the optimisation reaches good optimal set-points.

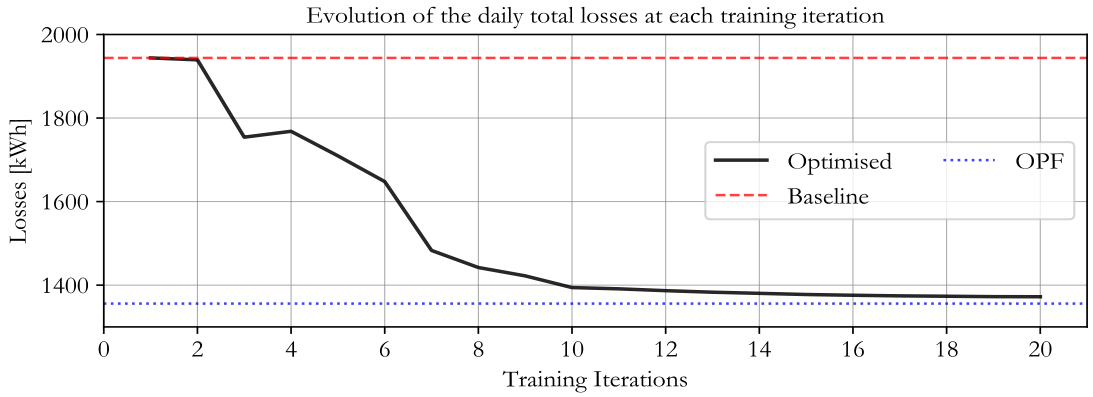



Figure 5.31: Total energy losses at every training iteration. The baseline is shown by the red dashed line whereas the minimum reached by the OPF is shown by the blue dotted line.

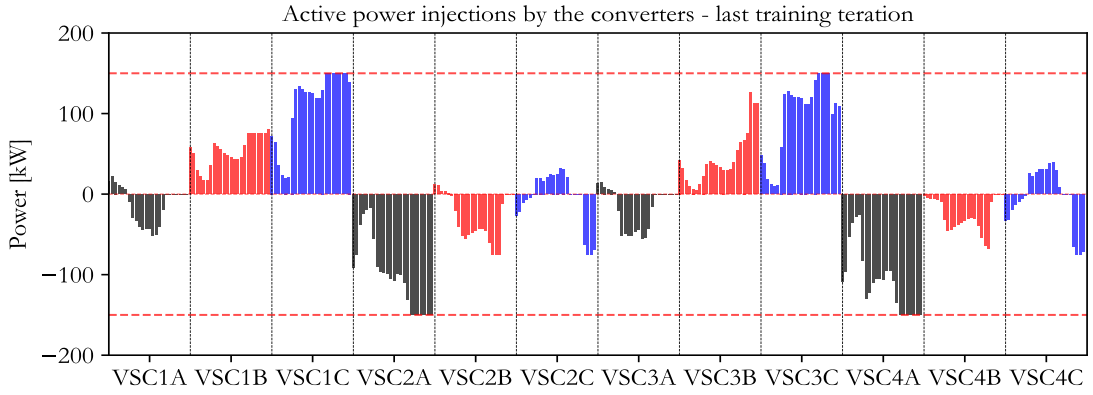



Figure 5.32: Active power injected or absorbed by converters set out sequentially for the 24hr period. The different colour of the bars indicate the phases. The red dashed line indicates the maximum active power that the converters can inject. The optimal power injections are obtained by the neural network (actor).

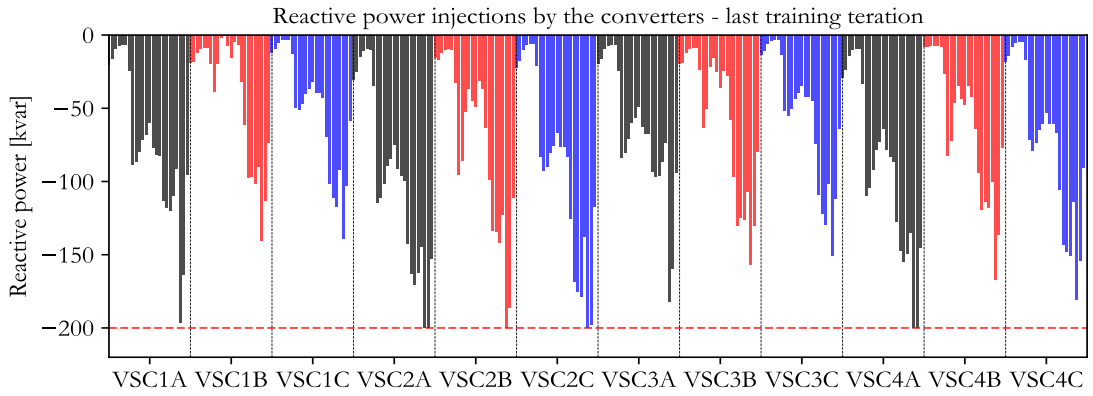



Figure 5.33: Reactive power injected or absorbed by converters set out sequentially for the 24hr period. The different colour of the bars indicate the phases. The red dashed line indicates the maximum active power that the converters can inject. The optimal power injections are obtained by the neural network (actor).

The effect of these power injections can be observed in Fig. 5.34, which compares the current supplied by the substation in the base case and with the effect of the converters. Similarly, a small performance gap is

observed, compared to Fig. 5.26, as the currents are not completely balanced. However, the goal of reducing the high phase imbalance, is achieved.

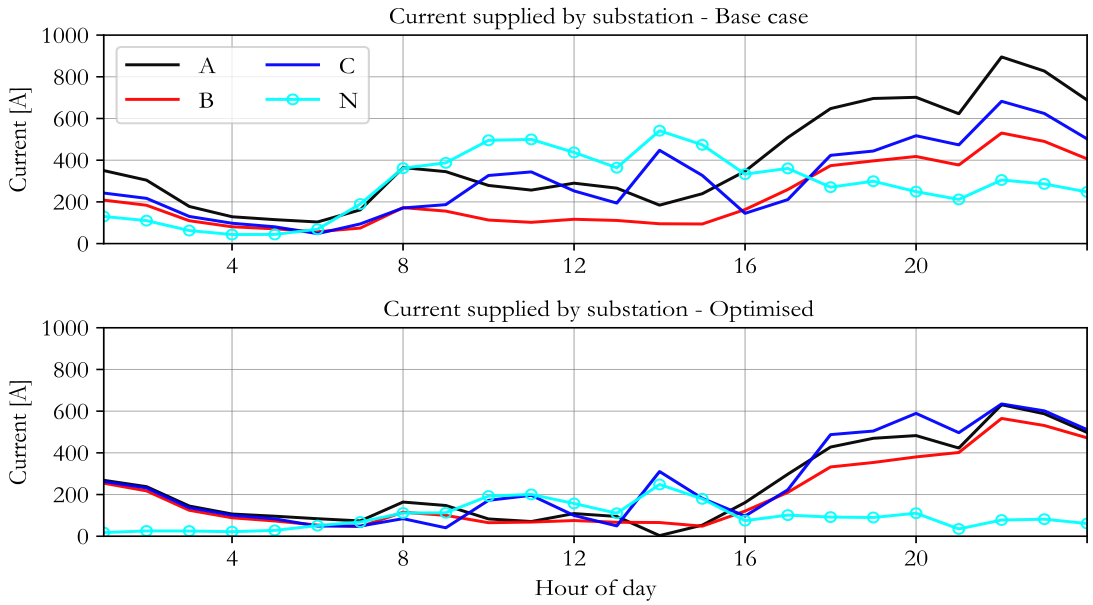



Figure 5.34: Comparison of the current supplied by the substation, for the baseline and for the optimised case.

The runtime for this case is summarised in Fig. 5.35. Due to the extra number of training iterations (TrainingIters) and the larger number of computations required to calculate the gradients and update the weights of the neural network, the optimisation runtime increased to 55 minutes from 35 minutes in the previous case (Fig. 5.29). The overhead runtime of building the computational graph and setting up the gradient operations remains low compare to the training. Once trained, the neural network can recall the optimal actions fast enough for real-time (even faster than a power flow calculation) taking approximately 10 ms per timestep and uses only the available measurements as input.

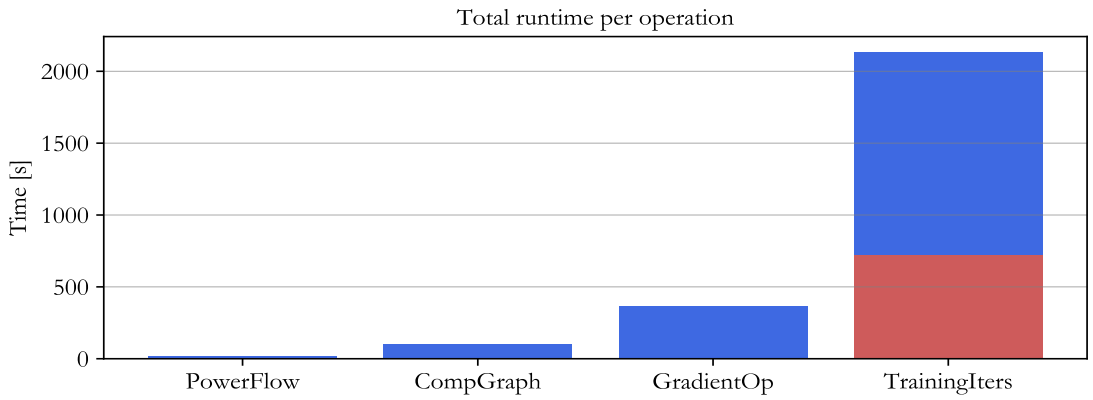



Figure 5.35: Runtime for each operation. For the training iterations, the red bar indicates the time taken by the first training iteration, the blue bar indicates the time taken by the subsequent 24 iterations.

Table 5.6 compares the results and runtime of the model-based approach against the power flow and OPF. The offline and online runtime results are given in seconds for a single timestep to make a more direct comparison, whereas the total runtime is given in minutes for all 24 timesteps. The losses represent the daily totals.

Table 5.6: Results and runtime per-time step IEEE 123-bus test feeder.

| Case | Total energy losses [kWh] | Runtime offline [s] | Runtime online [s] | Total runtime [min] |
|---|---|---|---|---|
| Base power flow | 1943.89 | 0.047 | | 0.01708 |
| OPF | 1355.80 | 365 | | 146 |
| Model-based - no action network | 1359.96 | 1280 | 30 | 35 |
| Model-based - with action network | 1402.10 | 1500 | 0.010 | 55 |

## 5.4 Discussion

Reinforcement learning, like any optimisation method, will have advantages and disadvantages that make it suited to some problems and not others. The case studies presented in this chapter were intended to establish whether RL is a good candidate for ANM in distribution networks.

Model-free RL schemes have an improved computational efficiency compared to the OPF due to their two-stage format (offline training on historical data and online use of a trained neural network). However, they require large volumes of training data, in the form of interactions with the environment. Also, as it was described in section 5.2.2 and shown in the case study, model-free algorithms require an exploration strategy to successfully search the solution space. Thus the time required for training is significant. However, they have the inherent ability to overcome issues with partial observability, because they can learn useful control policies by interacting only with a *black-box* environment representing the LV network with whatever observation are expected in practice.

Model-based alternatives are expected to be more data efficient given the direct interaction with the equations describing the network. Additionally, because of the use of gradient information, they do not require an exploration strategy and they are able to achieve faster convergence. Also, the discretisation of the action space, as done for the model-free method, is no longer required, resulting in an improved solution (i.e. faster convergence and improved local minima). One of the disadvantages of model-based approach is that all the operations that form the system need to be stored in a computational graph and are thus stored in memory. In the case of large LV networks, this means storing the fixed-point iterations of the power flow algorithm (Algorithm 3) and the gradient operations. Thus, although the power flow runtime increases only slightly

with the size of the network, the gradient operation takes considerably more time (twice as much as setting up the computational graph as shown in Fig. 5.29), which might hinder the performance.

It is also important to note that similar to the OPF, global optimality cannot be guaranteed for the solutions obtained by both model-based and model-free algorithms. However, extensions to the gradient descent algorithm, such as gradient descent with momentum, adaptive learning rate and gradient normalisation, can help overcome this potential problem [166, 168]. These enhancements have been shown to perform well when training deep neural networks for control problems [176]. Additionally, because the weights of the neural network are initialised randomly, several searches in the solution space can be performed by using different initial conditions to discover good local optima.

In both model-free and model-based RL methods, the constraints associated with the network parameters and the converters cannot be formulated as straightforward as in the constrained optimisation formulation of the OPF. The model-free version relies on applying large penalties to those actions that result in a constraint violation or a termination of the episode. The model-based approach can guarantee that the fixed-point algorithm converges to a feasible solution within the set tolerance for the voltage vector, as in (2.28), and that the converters' power injections are within the limits. However other constraints are not straightforward to implement.

Although both the model-free and model based methods can generalise to changes in the loading conditions, this is subject to the amount of historical data available. The data should also be from representative loading conditions, for instance loading data from different seasons of the year would be require to ensure a good performance throughout the year. It's worth noting that if the physical configuration of the distribution network is changed (e.g. through reconfiguration) or the converters are replaced with larger rated converters, the trained model would no longer be valid. Thus, a new training round would be required to update its internal representation (i.e. the parameters of the neural network) and to ensure optimal behaviour. This is also the case with typical OPF schemes, which are also unable to adapt to such physical changes and require an updated model to calculate new optimal set-points. However, it has been demonstrated in other applications [212, 213] that learned representations embedded in neural networks from similar tasks can be transferred to different tasks to avoid training a new neural network from scratch. For example, a neural network trained to classify images of handwritten numbers can be re-purposed to classify images of handwritten letters, without training the whole network from scratch, by *freezing* some of its layers and only training the last layers. This is known as transfer learning [213] and in the case of a model trained for ANM, it could be used whenever physical changes occur to a particular LV network. This would give a significant advantage to both model-free and model-based methods versus the OPF.

## 5.5 Summary

This chapter started with a review of the literature on reinforcement learning techniques suitable for analysis and control of power systems. Reinforcement learning for the control of the power flows in LV networks was considered as an alternative to using an OPF method and as a robust alternative that can generalise, unlike supervised learning methods. In contrast to traditional nonlinear optimisation, this type of optimal control can use approximations to overcome the *curse of dimensionality* and to achieve the computational efficiency that would be needed for real-time operation. Unlike supervised learning, the training involves feedback from the system, either through a *black-box* representation or from an explicit model of the network through the power flow equations. In particular, a distinction was made between model-free and model-based methods and their advantages and disadvantages.

The first evaluation was of self-learning control based on deep Q networks (DQN) applied to control of a pair of power converters (a SOP). The goal was to equalise the loading across two LV substations throughout the day. It was shown that the control strategy was able to learn optimal control policies that maximise the cumulative reward over the episodes, based only on sparse measurements and feedback from the agent-environment interaction. This was achieved without prior knowledge of the LV network dynamics and without an explicit model. This model-free approach to RL shows good generalisation but some shortcomings were apparent, such as the need to discretize the action space to be able to use DQN as a learning algorithm. This made the training of the neural network slow (6 hours). Also, the results showed that this had a detrimental effect on the estimation of the optimal power injections (i.e. the optimal actions can only have a limited granularity) for the converters and on the objective function (i.e. a relaxed tolerance is used for the objective function). The error compared to the OPF was between 4-6%.

The second evaluation of RL was of a model-based control following the ADP framework. The intention was to overcome some of the limitations found with the model-free approach. The power flow equations, solved through a fixed-point algorithm, were represented as a computational graph to facilitate the calculation of the gradient of the cost function with respect to control actions via back-propagation. The set-points of the converters were optimised with gradient descent with the objective of minimising the active power losses.

The first model-based implementation acts in a similar manner as an OPF. It optimises the actions directly, once it knows all the parameters of the network then it search for the schedule of optimal actions that would result in the minimum losses. The method was compared to the results obtained by the OPF. The model-based method achieved a local minimum with less than 1% difference of the OPF, in a third of the runtime. Nonetheless, this method is difficult to implement in real-time operation without having complete observability of the network.

To overcome the challenge of limited measurements available, function approximators are required to bridge the gap between the optimal actions and the observable measurements. In the second model-based implementation, a neural network (actor) was trained to choose the actions rather than optimising the actions directly. In the same way as with the model-free approach, the neural network is trained only on partial measurements, those that are normally available to network operator. Although this was shown to decrease the accuracy of the algorithm (the local minimum reached was 3% from one obtained by the OPF), it makes it suitable for implementation in realistic settings.

Due to the use of the power flow equations and their gradient, both model-based implementations were shown to be computationally efficient and could be readily implemented in a server for real-time operation.

From the results presented in this chapter, it can be concluded that a departure from the traditional optimisation is needed to formulate controls that act optimally with the computational and observability requirements for real-time operation. For this, alternatives arise from two RL frameworks. Model-free approaches are easier to incorporate with existing measurement infrastructure as they do not require an explicit model of the network. However, they need a suitable design of the reward function and they need an exploration strategy so that they can learn appropriate actions and eventually converge to an optimal solution. Model-based approaches require a model of the network to be constructed. For this, the topology of the network is needed, together with data on the network components. Although network connectivity models are not widely available for distribution networks, recent work and trials have been conducted to improve the visibility of these networks as the DNOs transition towards DSOs [55, 158]. However, once developed a model-based algorithm can provide optimal solutions in a much shorter time compared to both the OPF and model-free strategies. Because of the useful gradient information, the solutions obtained by a model-based algorithm are more accurate than those of model-free strategies and are comparable to those of the OPF. Moreover, when used to train a neural network, the neural network can be used independently to produce the optimal set-points, as it stores the relationship between the power losses and the power set-points in its weights.

A summary of the main findings, together with the conclusions and a discussion on the future research directions, are provided in the next chapter.

# Chapter 6

# Conclusions

## 6.1 Review of Problem Statement

Distribution networks, and in particular LV networks, have historically been designed to be passive networks, where power is delivered from substations to consumers via long feeders often arranged and operated radially. Also, DNOs usually considered only worst case scenarios for planning purposes, the case with maximum or minimum demand or generation. Thus, it was not necessary, nor cost-effective, to monitor the individual lines and nodes; measurements of aggregate loading were sufficient to determine the maximum/minimum demand expected and the headroom needed. Because of this planning philosophy, active control of the network was not required. Hence, the availability of instrumentation and data acquisition present in LV networks is low, compared to other parts of the power system. While the topology and structure of distribution networks have not changed significantly over the years, the type of loads connected to them has changed dramatically in recent years. Driven by the efforts to move to a carbon-neutral economy, large volumes of electric vehicles and charging points are being connected to distribution networks, as was remarked in section 1.1 and illustrated in Fig 1.2. Paired with the inclusion of distributed generation, notably solar photovoltaic, and energy storage devices, this has resulted in a number of challenges for network operators.

Large variations in power flow, and even reversals of power flow, caused by these LCTs, put additional pressure on DNOs, forcing them to adopt a more proactive role when operating these networks. Some of the main challenges faced by network operators are: (1) new connections to the network need to be provided at the lowest cost to customers and within reasonable timescales, (2) major upgrades and reinforcements are costly and can result in supply interruption and (3) network upgrades have to be made proactively while keeping the power supply and avoiding customer disconnections. This creates an incentive for DNOs to keep existing infrastructure but to use it more efficiently. As presented in section 1.3, early signs of the need

for flexibility and potential solutions are already tangible, in the form of DSR and generation-management mechanisms. However, to fully unlock the latent network capacity in existing infrastructure, ANM schemes that allow a complete control over the power flows are required. In order to defer network reinforcement, improved asset utilisation is necessary.

Although traditional network reconfiguration can be used to provide an extra degree of freedom by meshing LV feeders through switching, it was shown in **Chapter 3** that its benefits are limited. The power flow depends on characteristics such as the network impedance, and the location of LCTs. One of the case studies showed an example where ANM through switching has an adverse impact and does not improve the network capacity. On the other hand, an ANM solution based on power converters offers full control of the power flows within the network, thus providing more flexibility. Installing power converters in LV networks can provide alternative power flow paths between substations, feeders and even between phases. This can shift power flows from points with a high utilisation of network assets to less loaded parts of the network. Power can be distributed more evenly between substations and feeders with complimentary load/generation profiles. This improves asset utilisation and generally enhances network operation, improving the voltage profile and reducing power losses.

However, devising controllers that perform well for this type of ANM can be challenging for a number of reasons. First, unlike ANM through switching, controlling the power flows requires a constant calculation of the optimal set-points. Second, the assumption that the DNO has full-observability of the network does not hold for LV networks due to the lack of monitoring infrastructure. Third, the large number of nodes and the inherent phase imbalance makes modelling these networks more difficult and computationally intensive compared to other parts of the system, especially when coupled with the variability of LCTs. As was discussed in **Chapter 3**, network operators have increased the number of available monitoring devices in recent years, particularly in LV substations. Hence, it is likely that in the near future some degree of observability will be available for the implementation of ANM schemes. For a controller to be viable for ANM, there are several characteristics it requires. It needs to: (1) incorporate a realistic model of the network, representing its multi-phase, imbalanced nature, (2) be computationally efficient to be employed in real-time control, (3) be able to operate under partial observability by making approximations and (4) be able to operate autonomously and adapt to changes in the loading conditions.

The objectives set for the research work reported in this thesis were: (1) to investigate whether power electronic converters could be used to improve the operational outcomes of distribution networks and increase their capacity for hosting LCTs, (2) to analyse the benefits and limitations of ANM control methods based on traditional optimisation, (3) to devise new control schemes capable of: a) achieving an operation of the network that is close to that obtained with off-line traditional optimisation, b) performing this whilst operating under partial observability and c) being computationally efficient for real-time operation.

## 6.2 Contributions

A review of the main characteristics of distribution networks and how they affect the assumptions that can be made when modelling these networks for power flow and OPF studies was done in **Chapter 2**. Power flow and OPF formulations, suitable for realistic distribution networks with a high phase imbalance were presented. An AC/DC OPF formulation, based on the current injection method, was developed in **Chapter 3** to analyse the effect of power converters embedded in distribution networks. This AC/DC OPF incorporates a model of the power converters, their associated power losses and their respective DC connections as part of an ANM scheme. The DC connections were modelled as a MTDC network, similar to the approach used for offshore connections via HVDC. However, unlike transmission applications, power flows between all three phases were modelled individually to cater for the high phase imbalance of LV networks. This gives extra flexibility to the converters, as they can exchange power from different phases to reduce phase imbalance. This was found to release extra latent capacity compared to only exchanging power from point to point. However, it also translated into an increased complexity for the optimisation compared to transmission networks.

A series of networks with different characteristics and configurations were analysed using the AC/DC OPF to obtain the optimal set points for the ANM scheme. The results showed that the AC/DC OPF can find control set-points that lead to an optimal network operating state based on a given objective function (e.g. losses minimisation, transformer equalisation). An examination of the AC/DC OPF runtime was carried out to assess its suitability for real-time control. The results showed that the increase in computational effort between a power flow and the OPF is approximately two orders of magnitude ( 1 second vs 2.5 hours, for 24 timesteps). Also, for both the power flow and the OPF, the computational time increases with the size of the network. The increased computational burden, together with the requirement for full observability of the network were seen to prevent the OPF application in real-time control. A review of the different alternatives proposed in previous literature to overcome the OPF limitations was carried out. Several approaches were reviewed, including linear approximations and convex relaxations. However, because of the nature of low-voltage networks, it was deemed necessary to keep a realistic representation of these networks.

To overcome the difficulty of long computation times of the OPF in an ANM context, an alternative method based on supervised learning, using neural networks, was proposed in **Chapter 4**. For this data-driven control, a deep neural network was trained using optimal set points obtained offline through non-linear optimisation. Despite shifting the computational burden to an offline training stage, this method was found to have some limitations. Similar to other supervised learning algorithms, it lacked the ability to generalise to states outside of the training set. Moreover, the accuracy of the trained model was not entirely acceptable, with an average error of between 5% and 12% depending on the network. Also, because of

the characteristics inherent to regression models, the handling of constraints, such as the converter power-balance, was limited. That is, the regression estimated the power for each converter individually and a post-processing stage was needed to ensure the power transferred from one side matched the other side.

In **Chapter 5**, a novel modelling framework based on RL was proposed. First, a control based on model-free RL was proposed in section 5.2 to overcome the limitations of supervised learning and to approach a form of generalisation. In this method, the optimal distribution of the power was modelled as a sequential decision making process. Deep Q learning was used to overcome the limitations of tabular Q-learning. The algorithm trained a neural network by interacting with a model of the LV network via a power flow engine and by receiving constant feedback on the impact of the resulting power injections. Because this approach uses a *black-box* training environment, and lacks an explicit model of the LV network, the computational burden of the initial training stage was large, taking several hours to train. However once trained, the model stored, in the weights of the neural network, a representation of the LV network response to the converter injections. This response comprises parameters such as the power losses, loading of the transformers and the phase imbalance, and can be regarded as a measure of the power flow dynamics. The optimal set-points produced by this control had a performance comparable to the traditional optimisation and it displayed an improved ability to generalise to changes in the loading conditions, compared to the regression model presented in section 4.3.1.

Finally, in section 5.3, a more robust method was proposed. This model-based RL control incorporated an explicit model of the network, through the power flow equations, into the optimisation. This way, the effect of the power injections on the LV network was modelled via the power flow equations directly rather than via a *black-box* training environment. This method takes advantage of the computational efficiency of the fixed-point algorithm for power flow calculations (with an average runtime of 50ms) and frames it as a computational graph. It then uses backpropagation and gradient descent to efficiently calculate the gradients and optimise the actions based on them.

On a first instance, the model-based approach can optimise the power injections directly through back-propagation and gradient descent. However it can also be extended to deal with the challenge of partial observability by training a neural network to produce the optimal set-points. The weights of the neural network are tuned via backpropagation, using the gradient information. During the training stage, only some of the measurements are made available as input to the neural network. Thus, once trained the neural network can produce optimal set-points using only those available measurements.

Both model-based implementations were able to produce accurate set-points for the converters, with errors of only 1% and 3% respectively, as shown in Table 5.6. Because of the use of gradient-based optimisation and automatic differentiation, they display a superior computational tractability compared to both traditional optimisation and model-free RL. The training time was between 35-55 minutes for an entire day of

operation, compared to more than 2 hours taken by the OPF and the several hours taken by the model-free algorithm. Furthermore, their computational efficiency could be further improved by using GPUs or TPUs, because of their computational graph structure. Also, because of the function-approximation properties of the neural network, which produces the optimised actions, the algorithm is able to generalise to changing loading conditions whilst keeping an acceptable performance compared to the OPF, which requires a full optimisation every time the loading changes.

## 6.3 Conclusions

From the literature reviewed in **Chapter 1** and **Chapter 3**, it is clear that as the adoption of LCTs accelerates, it will be hard for DNOs to reinforce the network at the same pace. Improved network flexibility is thus required to avoid reaching the thermal and voltage limits of the network components. Active management of the power flows, in combination with other complementary smart grid solutions, such as DSR, will become necessary to allow the installation of a higher number of LCTs. Using power converters to continuously manage the power flows in LV networks can alleviate the impacts of these technologies and allow a larger number of these connections. In particular, it has been shown in the case studies in section 3.4 that a lower phase imbalance can be achieved by using power converters, as part of an ANM scheme. Balancing the power flows between phases can reduce significantly the neutral current, as shown in the IEEE 123-bus test feeder, where a 70% decrease was achieved. This also results in reduced power losses, with an overall 30% reduction observed in the IEEE 123-bus test feeder. Also, as shown in the Lotherton and Canefields case study, peak loading for both transformers and feeders can be reduced by moving power flows from highly loaded components to less utilised parts of the network. As a consequence of this, the voltage profile is improved, the overall power losses are reduced, and latent network capacity is released. ANM schemes based on power converters can release a higher amount of latent capacity compared to ANM through switching, as illustrated by the total kVA capacity released (10-30 kVA) during peak hours for the Lotherton and Canefields substations. However, to implement an ANM scheme based on power converters requires a more sophisticated control compared to switching, and it requires a certain degree of observability in the network.

To enable the active management of the power flows, various control methods have been proposed, analysed and compared in this thesis. Evidence of the advantages and limitations of each of the proposed control strategies was provided through a series of case studies using realistic network models and demand profiles. The results from the case studies showed that although the OPF can produce the optimal set points for an ANM scheme with power converters and DC links, this optimisation approach is not practical for real-time implementation. To implement an ANM scheme using the OPF as a controller would require a new optimisation step for every new loading condition. Given its high computational runtime, in the range

of hours, the OPF is not practical when dealing with a high penetration of LCTs and thus is not suitable for real-time operation.

Adopting a supervised learning algorithm, such as regression, to reduce the computational burden of the OPF seems like the immediate choice. Although the accuracy of a regression algorithm is acceptable (approximately 5% for the IEEE 123-bus case and 12% for the European LV case) the training time is considerable and the OPF is still required to produce the training data for several scenarios. Moreover, a regression approach does not incorporate any feedback from the network model and thus cannot generalise to scenarios outside of the range of the training set. Because of the disadvantages of supervised learning, it was deemed required to propose a radical shift into how the optimisation is performed. Reinforcement learning offers an alternative framework for developing an optimal control strategy that maintains a realistic representation of LV networks whilst still enabling optimal behaviour and taking advantage of the function approximation properties of neural networks. Both model-free and model-based schemes using neural networks were proposed and compared to traditional optimisation.

For those applications where the topology of the network is not known or a network model is difficult to construct, the model-free approach offers the best alternative, as the neural network can be trained directly, using only measurements (e.g., from the monitoring available in the OpenLV project [119]). This model-free approach produces setpoints with an average error of only 4% and 6%, with respect to the OPF. Moreover, this approach was able to generalise, producing optimal set-points, which resulted in loading equalisation between the transformers, for a different day (Monday) that the one trained on (Sunday), with a RMSE of 2.682 kW (less than 10% of the rating of the converter)). Nonetheless, a careful design of the reward function and the termination episodes is required (as described in 5.2.2). Additionally, the training needs to be done offline, as the time taken is considerable (6 hours for the case study in section 5.2.3).

In the case of the model-based approach, two main advantages are a fast training stage and a high optimisation accuracy. This approach achieved a local minimum with less than 1% difference compared to the OPF and it had a total runtime between 60-70% less than the time taken by the OPF, as shown in Table 5.6. However, a disadvantage compared to the model-free approach is that a full model of the LV network is required to carry out the training. This means that the topology and connectivity of the network are needed, together with historical data. Although detailed models of distribution networks are not typically available to DNOs, recent projects have developed such models from geographical information [159, 214]. Hence, a model of the Lotherton and Canefields substations could be constructed from the geographical information available to the DNO as shown in Fig. 3.29 and Fig. 3.31. This model could then be used to train a neural network to make optimal decisions based on the available measurements, such as those from the OpenLV project [119]. Table 6.1 provides a summary comparing the characteristics of the different approaches analysed throughout the thesis.

To further integrate LCTs into distribution networks requires ANM schemes that can deal with the great variability of LCTs, which can overcome the uncertainty arising from partial observability and that can adapt as the loading conditions change. Hence, control algorithms offering a compromise between accuracy, observability requirements and computational tractability are needed. For a truly autonomous scheme that meets the requirements to be implemented in real-time operation, a combination of RL algorithms that incorporate feedback from the network, together with function approximators, is essential to overcome the *curse of dimensionality* and to approach a form of generalisation.

This thesis provided one of the first assessment of the benefits of formulating the optimal power flow problem using RL methods and neural networks. It was demonstrated in the case studies that RL represents a strong alternative to the traditional OPF. Unlike supervised learning, a controller based on RL and neural networks can learn to produce optimal setpoints by interacting with the power flow equations directly or with a *black-box* training environment, and thus can generalise to changes in the loading conditions. The results from this research will serve as a base for future implementations of the OPF that are computationally efficient and can be scaled for use in large distribution networks when operating in real-time. The findings will be of interest to DNOs when considering the implementation of ANM schemes.

Table 6.1: Comparison between alternative methods proposed and traditional optimisation.

| | OPF | Supervised learning | RL model-free | RL model-based |
|---|---|---|---|---|
| **Optimisation framework** | Nonlinear constrained optimisation using a method such as the interior point method. | Needs an optimisation engine (e.g., OPF) to produce a dataset of optimised setpoints. A regression algorithm can then be trained on the dataset. | A neural network is trained using constant interactions with a power flow engine, using a reward function and a memory buffer to store past experiences. | Gradient-based optimisation of the power flow equations, framed as a computational graph. Simultaneous training of a neural network. |
| **Calculation of derivatives** | Derivatives are calculated using numerical differentiation. | For the case of a neural network, the derivatives are calculated using automatic differentiation. | Derivatives are calculated using automatic differentiation. | Derivatives are calculated using automatic differentiation. |
| **Global optimality** | Cannot guarantee global optimality. However, convex relaxation techniques can be used to search for global optimum. | The optimal setpoints are only as good as those obtained by the optimisation used to produce the training dataset. Cannot guarantee global optimality | Cannot guarantee global optimality but modern heuristic techniques can be used to search solution space (momentum, adaptive learning rate, etc.). | Cannot guarantee global optimality but modern heuristic techniques can be used to search solution space (momentum, adaptive learning rate, etc.). |
| **Observability requirements** | Needs full visibility at all times to produce optimal setpoints. | Needs full visibility only during training. Once trained, the neural network can operate under partial observability. | Does not need full visibility. The neural network can be trained and operate on available measurements. | Needs full visibility only during training. Once trained, the neural network can operate under partial observability. |
| **Computational efficiency** | Large computational burden of nonlinear, non-convex optimisation. Total runtime in the range of hours. | Two-stage format. The dataset and the training are performed offline with a runtime in the range of hours. The neural network online recall is very fast. | Two-stage format. The training stage requires an exploration strategy and has a runtime in the range of hours. The neural network online recall is very fast. | Two-stage format. The training stage does not require exploration and has a runtime in the range of minutes. The neural network online recall is very fast. |
| **Generalisation ability** | A new optimisation is needed for every timestep regardless of similarity between loading conditions. | The regression algorithm cannot extrapolate to loading conditions outside of range of the training dataset. | The constant feedback and the reward function allows the neural network to generalise to out-of-sample loading conditions. | Feedback from the gradient information allows the neural network to generalise to out-of-sample loading conditions. |

## 6.4 Future Research

Designing controllers for ANM schemes using power converters has several design and operation challenges that have been discussed throughout the thesis. However, several questions still need to be resolved before real-time operation can be implemented. Those areas that require further research will be discussed in this section.

It is recommended that further studies are carried out on a larger number of networks to validate the performance and scalability of the methods presented. For instance, the model-free approach should be implemented with a network such as the IEEE 123-bus and the model-based approach should be tested using the IEEE 8500-node test feeder. Also, in some case studies, average load and generation profiles were used for all the nodes, due to the lack of smart meter data. To improve this, a further study could assess the effect of having realistic individual profiles for each node and the impact of this additional diversity on the control scheme. Another approach that could be examined in future research, to overcome the lack of loading data and to increase the number of scenarios used for training, would be the use of a probabilistic method. For instance, the Monte Carlo method could be combined with the RL algorithms proposed to cover a broader range of possible loading scenarios. This way, the loading/generation information would be sampled from a probabilistic distribution rather than being an average of historical data.

This thesis focused on the high-level control of power electronic converters for use in ANM schemes under steady-state conditions. For this, one of the main assumptions made was that the power set-points produced by this control can be transmitted to the converters' lower levels of control in a straightforward way and with the update rate needed to cope with changes in loading conditions. However, the internal dynamics of the converters might require further analysis to validate this assumption. The analysis of the control scheme under faults was not analysed but is suggested for future work. In particular, the methods based on RL, which require offline training, would need to be combined with lower-level controllers. Designing a coordinated hierarchical control might be required, where the optimal set-points are calculated by the high-level control based on the overall network conditions (e.g., using a scheme based on RL) and a low-level control, based on traditional algorithms (e.g., the one proposed in [129]), which produces the setpoints for the converter current injections using local measurements. The low-level controller would also ensure the power converter is protected against faults and would avoid internal damage of the converter in situations where the global control does not have sufficient information to make an optimal decision. Further work and the implementation of both the high-level and the low-level controls in a hardware test-bed is required to successfully validate the performance. Coordination with the traditional protection systems also needs to be analysed in more detail to guarantee a secure operation of the ANM scheme under faults.

Although some of the control methods proposed in this study can deal with the partial observability

found in LV networks, via function approximators, their implementation still follows a centralised approach. Because of this, suitable communication infrastructure is necessary to transmit the power set-points to the converters whenever they are located downstream along the feeders. Further work is needed to adapt its performance in a distributed or multi-agent manner, to improve its resilience to communication failures. The latest developments in RL [177] have shown that learning methods can be adapted in a multi-agent fashion. Thus, it is recommended that further research be undertaken in this area to adapt the controls presented here to allow for interactions between multiple converters. Also, the advantages and suitability of transfer learning would be worth exploring to avoid training neural networks from scratch.

Another aspect worth investigating in more detail is the potential for including online learning as part of the control strategy. In this case, a neural network would be trained offline using RL and representative historical data or a probabilistic model. However, once trained, rather than only producing the setpoints at every timestep, it could also learn from the real-time interactions, as proposed in [215] and more recently in [216].

One of the assumptions made through the work was that a model of the distribution network was available or could be built. However, at present this might be difficult to do for some LV networks. Moreover, the model of the network would need to incorporate the different running arrangements used for contingencies. Although building a model of the distribution network could be time consuming for DNOs, the benefits of this are significant and have been demonstrated in the case studies. Future work should explore the development of LV models and integrated HV/LV models with the objective of trailing ANM schemes.

Lastly, this thesis focused on the benefits to the network infrastructure from the technical point of view. However, an economic evaluation of the viability and the monetary advantages of the ANM schemes proposed is suggested. This could help DNOs make informed decisions on the long-term cost efficiency of the ANM scheme proposed compared to traditional network reinforcement.

# Bibliography


[1] IEA, "Global EV outlook 2020," International Energy Agency, Tech. Rep., June 2020.

[2] N. Kittner, F. Lill, and D. M. Kammen, "Energy storage deployment and innovation for the clean energy transition," *Nature Energy*, vol. 2, no. 9, p. 17125, 2017. [Online]. Available: https://doi.org/10.1038/nenergy.2017.125

[3] B. Nykvist, F. Sprei, and M. Nilsson, "Assessing the progress toward lower priced long range battery electric vehicles," *Energy Policy*, vol. 124, pp. 144 – 155, 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0301421518306487

[4] National Grid, "UK Future Energy Scenarios," National Grid, Report, July 2013 2013.

[5] C. Barteczko-Hibbert, "Customer-led network revolution: After diversity maximum demand (ADMD) report," Tech. Rep.

[6] Mingyang Sun, I. Konstantelos, and G. Strbac, "Analysis of diversified residential demand in london using smart meter and demographic data," in *2016 IEEE Power and Energy Society General Meeting (PESGM)*, 2016, pp. 1–5.

[7] D. H. O. McQueen, P. R. Hyland, and S. J. Watson, "Monte carlo simulation of residential electricity demand for forecasting maximum demand on distribution networks," *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1685–1689, 2004.

[8] SAE, "SAE electric vehicle and plug in hybrid electric vehicle conductive charge coupler - J1772-201710," 2017.

[9] UK EVSE, "Making the right connections: General procurement guidance for electric vehicle charge points," UK Electric Vehicle Supply Equipment Association, Report, 2015. [Online]. Available: http://ukevse.org.uk/content/uploads/2015/04/CNX2849-UK-EVSE-Procurement-Guide-AW-WEB.pdf

[10] BNEF, "New energy outlook 2016," BNEF, Report, 2016.

[11] G. J. Offer, "Automated vehicles and electrification of transport," *Energy & Environmental Science*, vol. 8, no. 1, pp. 26–30, 2015. [Online]. Available: http://dx.doi.org/10.1039/C4EE02229G

[12] J. Arbib and T. Seba, "Rethinking transportation 2020-2030: The disruption of transportation and the collapse of the internal-combustion vehicle and oil industries," RethinkX, Report, May 2017 2017.

[13] M. Neaimeh, S. D. Salisbury, G. A. Hill *et al.*, "Analysing the usage and evidencing the importance of fast chargers for the adoption of battery electric vehicles," *Energy Policy*, vol. 108, pp. 474 – 486, 2017. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0301421517303877

[14] Open Charge Map, "Open Charge Map," 2020, [Accessed 04-12-2020]. [Online]. Available: https://openchargemap.org/site/

[15] Energy Technologies Institute, "Smart systems and heat," Energy Technologies Institute, Report, 2015.

[16] Element Energy & E4tech, "Cost analysis of future heat infrastructure options," Element Energy & E4tech, Report, 2018.

[17] Element Energy, "Distribution future energy scenarios UK Power Networks," Element Energy, Report, 2020.

[18] S. Watson, K. Lomas, and R. Buswell, "Decarbonising domestic heating: What is the peak gb demand?" *Energy Policy*, vol. 126, pp. 533 – 544, 2019. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0301421518307249

[19] BEIS, "Solar photovoltaics deployment," 27 February 2020 2020. [Online]. Available: https://www.gov.uk/government/statistics/solar-photovoltaics-deployment

[20] IRENA, "Future of solar photovoltaic: Deployment, investment, technology, grid integration and socio-economic aspects," IRENA, Report, November 2019.

[21] T. Braunholtz-Speight, M. Sharmina, E. Manderson *et al.*, "Business models and financial characteristics of community energy in the uk," *Nature Energy*, p. 1–9, 2020.

[22] UK Power Networks, "Validation of Photovoltaic (PV) Connection Assessment Tool - Closedown Report," UKPN, Report, March 2015.

[23] Y. Wang, N. Zhang, Q. Chen *et al.*, "Data-Driven Probabilistic Net Load Forecasting With High Penetration of Behind-the-Meter PV," *IEEE Transactions on Power Systems*, vol. 33, no. 3, pp. 3255–3264, 2018.

[24] CPUC, "Beyond 33% renewables: Grid integration policy for a low-carbon future," California Public Utilities Commission, Report, 11/25/2015 2015.

[25] D. Della Giustina, M. Pau, P. A. Pegoraro *et al.*, "Electrical distribution system state estimation: measurement issues and challenges," *IEEE Instrumentation Measurement Magazine*, vol. 17, no. 6, pp. 36–42, December 2014.

[26] UK Power Networks, "Application to Innovation Roll-Out Mechanism: LV network visibility and control," UKPN, Report, 2017. [Online]. Available: https://www.ofgem.gov.uk/ofgem-publications/117914

[27] A. J. Roscoe and G. Ault, "Supporting high penetrations of renewable generation via implementation of real-time electricity pricing and demand response," *IET Renewable Power Generation*, vol. 4, no. 4, pp. 369–382, 2010.

[28] A. Ballanti and L. F. Ochoa, "Voltage-led load management in whole distribution networks," *IEEE Transactions on Power Systems*, vol. 33, no. 2, pp. 1544–1554, 2018.

[29] ELEXON, "Load profiles and their use in electricity settlement," ELEXON, Report, October 2018. [Online]. Available: https://www.elexon.co.uk/documents/training-guidance/bsc-guidance-notes/load-profiles/

[30] R. Silversides, T. Luth, and T. C. Green, "Power electronics in distribution system management," HubNet, Report, 2015. [Online]. Available: http://www.hubnet.org.uk/filebyid/633/PE_Distribution.pdf

[31] A. Navarro-Espinosa and L. F. Ochoa, "Probabilistic impact assessment of low carbon technologies in LV distribution systems," *IEEE Transactions on Power Systems*, vol. 31, no. 3, pp. 2192–2203, 2016.

[32] J. Wamburu, S. Lee, P. Shenoy, and D. Irwin, "Analyzing distribution transformers at city scale and the impact of EVs and storage," in *Proceedings of the Ninth International Conference on Future Energy Systems*, ser. e-Energy '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 157–167. [Online]. Available: https://doi.org/10.1145/3208903.3208925

[33] A. R. Ahmadi, T. Manandhar, J. Barros *et al.*, "UK Power Networks' experience of managing flexible distributed generation from planning to operation," *CIRED - Open Access Proceedings Journal*, vol. 2017, no. 1, pp. 2032–2036, 2017.

[34] UK Power Networks, "B2 impact of electric vehicles and heat pump loads on network demand profiles," UK Power Networks, Report, September 2014.

[35] Energy Networks Association, "Mapping current SO, TO and DNO processes," ENA, Report, June 2017.

[36] Ofgem, "Electricity distribution price control review final proposals - allowed revenue - cost assessment appendix," Report, December 2009.

[37] Scottish and Southern Energy Power Distribution, "Climate change adaptation report," Scottish and Southern Energy Power Distribution, Report, July 2015.

[38] Energy Networks Association, "Distribution charges overview," 2019, [Online; accessed 6-February-2020]. [Online]. Available: http://www.energynetworks.org/electricity/regulation/distribution-charging/distribution-charges-overview.html

[39] W. H. Kersting, *Distribution System Modeling and Analysis*, 3rd ed. Boca Raton, Florida: CRC Press, 2012.

[40] H. Pezeshki, P. J. Wolfs, and G. Ledwich, "Impact of high PV penetration on distribution transformer insulation life," *IEEE Transactions on Power Delivery*, vol. 29, no. 3, pp. 1212–1220, 2014.

[41] P. Bryan, R. Broderick, B. Mather *et al.*, "On the path to sunshot: Emerging issues and challenges in integrating solar with the distribution system," NREL, Report, 2016.

[42] M. Coddington, B. Mather, B. Kroposki *et al.*, "Updating interconnection screens for pv system integration," NREL, Report, 2012.

[43] EPRI, "Alternatives to the 15% rule: Final project summary," EPRI, Report 3002006594, 2015.

[44] A. Ballanti and L. F. Ochoa, "On the integrated pv hosting capacity of mv and lv distribution networks," in *2015 IEEE PES Innovative Smart Grid Technologies Latin America (ISGT LATAM)*, 2015, pp. 366–370.

[45] J. Coignard, P. MacDougall, F. Stadtmueller, and E. Vrettos, "Will electric vehicles drive distribution grid upgrades?: The case of California," *IEEE Electrification Magazine*, vol. 7, no. 2, pp. 46–56, June 2019.

[46] J. Soares, H. Morais, T. Sousa *et al.*, "Day-ahead resource scheduling including demand response for electric vehicles," *IEEE Transactions on Smart Grid*, vol. 4, no. 1, pp. 596–605, 2013.

[47] P. Papadopoulos, S. Skarvelis-Kazakos, I. Grau *et al.*, "Electric vehicles' impact on British distribution networks," *IET Electrical Systems in Transportation*, vol. 2, pp. 91–102(11), September 2012. [Online]. Available: https://digital-library.theiet.org/content/journals/10.1049/iet-est.2011.0023

[48] S. Bruno, S. Lamonaca, G. Rotondo *et al.*, "Unbalanced three-phase optimal power flow for smart grids," *IEEE Transactions on Industrial Electronics*, vol. 58, no. 10, pp. 4504–4513, 2011.

[49] E. Dall'Anese, H. Zhu, and G. B. Giannakis, "Distributed optimal power flow for smart microgrids," *IEEE Transactions on Smart Grid*, vol. 4, no. 3, pp. 1464–1475, 2013.

[50] F. Shahnia, A. Ghosh, G. Ledwich, and F. Zare, "Predicting voltage unbalance impacts of plug-in electric vehicles penetration in residential low-voltage distribution networks," *Electric Power Components and Systems*, vol. 41, no. 16, pp. 1594–1616, 2013. [Online]. Available: http://dx.doi.org/10.1080/15325008.2013.834004

[51] N. Leemput, F. Geth, J. V. Roy *et al.*, "Impact of electric vehicle on-board single-phase charging strategies on a flemish residential grid," *IEEE Transactions on Smart Grid*, vol. 5, no. 4, pp. 1815–1822, 2014.

[52] A. Dubey and S. Santoso, "Electric vehicle charging on residential distribution systems: Impacts and mitigations," *IEEE Access*, vol. 3, pp. 1871–1893, 2015.

[53] J. S. Marshall, P. D. Hines, J. D. Zhang *et al.*, "Modeling the impact of electric vehicle charging on heat transfer around underground cables," *Electric Power Systems Research*, vol. 97, pp. 76 – 83, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378779612003586

[54] L. McDonald, A. R. Ahmadi, S. Do, and S. Georgiopoulos, "Enabling distributed energy resources to enter the energy market and supporting the evolution to a distribution system operator," *CIRED - Open Access Proceedings Journal*, vol. 2017, no. 1, pp. 2715–2718, 2017.

[55] Western Power Distribution, "SDRC 1 - Detailed design of the overall OpenLV solution," Report, October 2017. [Online]. Available: https://openlv.net/resources/project-information/

[56] S. E. Networks, "Lv engine project progress report 2019," SP Energy Networks, Report, 2020.

[57] G. Celli, F. Pilo, G. Pisano *et al.*, "Meshed vs. radial mv distribution network in presence of large amount of DG," in *IEEE PES Power Systems Conference and Exposition.*, 2004, Conference Proceedings, pp. 709–714 vol.2.

[58] A. Navarro-Espinosa, L. F. Ochoa, and D. Randles, "Assessing the benefits of meshed operation of LV feeders with low carbon technologies," in *ISGT*, 2014, Conference Proceedings, pp. 1–5.

[59] UK Power Networks, "Flexible urban networks – selection criteria and approach (sdrc 9.1)," UK Power Networks, Report, 2015.

[60] M. Loos, S. Werben, and J. C. Maun, "Circulating currents in closed loop structure, a new problematic in distribution networks," in *2012 IEEE Power and Energy Society General Meeting*, July 2012, pp. 1–7.

[61] C. Long, J. Wu, L. Thomas, and N. Jenkins, "Optimal operation of soft open points in medium voltage electrical distribution networks with distributed generation," *Applied Energy*, vol. 184, pp. 427–437, 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0306261916314702

[62] UK Power Networks, "Flexible urban networks – low voltage successful demonstrations of enhanced modes of operation of power electronics devices (sdrc 9.5)," UK Power Networks, Report, 2015.

[63] Weedy, *Electric power systems*, 5th ed. Chichester, West Sussex, UK: John Wiley & Sons, Ltd., 2012.

[64] J. J. Grainger and W. D. Stevenson, *Power system analysis*. New York, NY: McGraw-Hill, 2016.

[65] BVG Associates Limited, "Guide to an offshore wind farm," BVG Associated Limited, Report, 2019.

[66] A. Gomez-Exposito, A. Conejo, and C. Canizares, *Electric Energy Systems*, 2nd ed. Boca Raton, Florida: CRC Press, 2018.

[67] E. Lakervi and E. J. Holmes, *Electricity Distribution Network Design*, 2nd ed. IET, 2003.

[68] UK Power Networks, "EHV Network Design - EDS 08-4000," UKPN, Report, 2017. [Online]. Available: http://library.ukpowernetworks.co.uk/library/asset/2192b07d-b35c-4f23-a5b8-c2b720c0776G/EDS+08-4000+EHV+Network+Design.pdf

[69] I. Hernando-Gil, B. Hayes, A. Collin, and S. Djokić, "Distribution network equivalents for reliability analysis. part 1: Aggregation methodology," in *IEEE PES ISGT Europe 2013*, Oct 2013, pp. 1–5.

[70] A. J. Urquhart and M. Thomson, "Series impedance of distribution cables with sector-shaped conductors," *IET Generation, Transmission Distribution*, vol. 9, no. 16, pp. 2679–2685, 2015.

[71] H. Ahmadi, J. R. Marti, and A. v. Meier, "A linear power flow formulation for three-phase distribution systems," *IEEE Transactions on Power Systems*, vol. 31, no. 6, pp. 5012–5021, 2016.

[72] B. S. Institution, "Lv and mv polymeric insulated cables for use by distribution and generation utilities part 3: Specification for distribution cables of rated voltage 0.6/1 kv," British Standards Institution, Standard BS 7870-3.50:2011, 2011. [Online]. Available: https://bsol.bsigroup.com/Bibliographic/BibliographicInfoData/000000000030262289

[73] Electricity North West Limited, "LV capacitor placement study phase 1 - capacitor placement and sizing," Report, 27 March 2013 2013.

[74] K. Miu and M. Kleinberg, "Impact studies of unbalanced multi-phase distribution system component models," in *IEEE PES General Meeting*, July 2010, pp. 1–4.

[75] S. Thomas, *Electric power distribution handbook*, 2nd ed. Boca Raton, Florida: CRC Press, 2014.

[76] CIGRE Task Force C6.04, "Benchmark systems for network integration of renewable and distributed energy resources." CIGRE, Report, April 2014 2014.

[77] M. J. Zangs, T. Yunusov, W. Holderbaum, and B. Potter, "On-line adjustment of battery schedules for supporting lv distribution network operation," in *2016 International Energy and Sustainability Conference (IESC)*, June 2016, pp. 1–6.

[78] W. H. Kersting, "A method to teach the design and operation of a distribution system," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-103, no. 7, pp. 1945–1952, July 1984.

[79] I. Hernando-Gil, H. Shi, F. Li *et al.*, "Evaluation of fault levels and power supply network impedances in 230/400 v 50 hz generic distribution systems," *IEEE Transactions on Power Delivery*, vol. 32, no. 2, pp. 768–777, April 2017.

[80] UK Power Networks, "Business Plan - Annex 7: Losses Strategy," Report, March 2014 2014. [Online]. Available: https://www.ukpowernetworks.co.uk/losses/static/pdfs/ukpn-business-plan-2015-2023.ded5a34.pdf

[81] UK Power Networks, "Loading of secondary distribution transformers," Report, 2015.

[82] A. M. Sasson and F. J. Jaimes, "Digital methods applied to power flow studies," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-86, no. 7, pp. 860–867, July 1967.

[83] W. F. Tinney and C. E. Hart, "Power flow solution by newton's method," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-86, no. 11, pp. 1449–1460, 1967.

[84] A. v. Meier, *Electric power systems a conceptual introduction*, ser. Wiley survival guides in engineering and science. Hoboken, N.J.: IEEE Press : Wiley-Interscience, 2006.

[85] I. Richardson, M. Thomson, D. Infield, and C. Clifford, "Domestic electricity use: A high-resolution energy demand model," *Energy and Buildings*, vol. 42, no. 10, pp. 1878 – 1887, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0378778810001854

[86] A. Navarro, L. F. Ochoa, and P. Mancarella, "Learning from residential load data: Impacts on LV network planning and operation," in *2012 Sixth IEEE/PES Transmission and Distribution: Latin America Conference and Exposition (T D-LA)*, Sep. 2012, pp. 1–8.

[87] B. Stephen, A. J. Mutanen, S. Galloway *et al.*, "Enhanced load profiling for residential network customers," *IEEE Transactions on Power Delivery*, vol. 29, no. 1, pp. 88–96, Feb 2014.

[88] C. L. Fortescue, "Method of symmetrical co-ordinates applied to the solution of polyphase networks," *Transactions of the American Institute of Electrical Engineers*, vol. XXXVII, no. 2, pp. 1027–1140, 1918.

[89] G. Chicco and A. Mazza, "100 years of symmetrical components," *Energies*, vol. 12, no. 3, p. 450, Jan 2019. [Online]. Available: http://dx.doi.org/10.3390/en12030450

[90] M. Thomson and D. G. Infield, "Network power-flow analysis for a high penetration of distributed generation," *IEEE Transactions on Power Systems*, vol. 22, no. 3, pp. 1157–1162, 2007.

[91] B. S. Institution, "Voltage characteristics of electricity supplied by public electricity networks," British Standards Institution, Standard BS EN 50160:2010+A3:2019, 2010. [Online]. Available: https://bsol.bsigroup.com/Bibliographic/BibliographicInfoData/000000000030365905

[92] P. Pillay and M. Manyage, "Definitions of voltage unbalance," *IEEE Power Engineering Review*, vol. 21, 2001.

[93] F. Ghassemi and M. Perry, "Review of voltage unbalance limit in the GB grid code CC.6.1.5 (b)," National Grid, Report, October 2014 2014.

[94] T. V. Cutsem and C. Vournas, *Voltage stability of electric power systems*. Springer Verlag, 1998.

[95] R. C. Dugan and T. E. McDermott, "An open source platform for collaborating on smart grid research," in *2011 IEEE Power and Energy Society General Meeting*, 2011, Conference Proceedings, pp. 1–7.

[96] R. F. Arritt and R. C. Dugan, "The ieee 8500-node test feeder," in *IEEE PES T D 2010*, 2010, pp. 1–6.

[97] H. W. Dommel and W. F. Tinney, "Optimal power flow solutions," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-87, no. 10, pp. 1866–1876, 1968.

[98] J. A. Momoh, *Electric Power System Applications of Optimization*, 2nd ed. Boca Raton, Florida: CRC Press, 2008.

[99] J. Zhu, *Optimization of Power System Operation*, 2nd ed. John Wiley & Sons, Inc, 2015.

[100] F. Capitanescu, M. Glavic, and L. Wehenkel, "An interior-point method based optimal power flow," June 2005.

[101] S. Frank and S. Rebennack, "An introduction to optimal power flow: Theory, formulation, and examples," *IIE Transactions*, vol. 48, no. 12, pp. 1172–1197, 2016. [Online]. Available: http://dx.doi.org/10.1080/0740817X.2016.1189626

[102] S. Frank, I. Steponavice, and S. Rebennack, "Optimal power flow: a bibliographic survey I," *Energy Systems*, vol. 3, no. 3, pp. 221–258, 2012. [Online]. Available: http://dx.doi.org/10.1007/s12667-012-0056-y

[103] J. Lavaei and S. H. Low, "Convexification of optimal power flow problem," in *2010 48th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, Sept 2010, pp. 223–232.

[104] J. Lavaei and S. H. Low, "Zero duality gap in optimal power flow problem," *IEEE Transactions on Power Systems*, vol. 27, no. 1, pp. 92–107, 2012.

[105] J. Gondzio, "Crash start of interior point methods," *European Journal of Operational Research*, vol. 255, no. 1, pp. 308 – 314, 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0377221716303502

[106] Ofgem, "Electricity distribution units and loss percentages summary - factsheet," Report, August 2010.

[107] G. Strbac, P. Djapic, D. Pudjianto *et al.*, "Strategies for reducing losses in distribution networks." Imperial College London - UKPN, Report, February 2018.

[108] S. Paudyal, C. A. Canizares, and K. Bhattacharya, "Optimal operation of distribution feeders in smart grids," *IEEE Transactions on Industrial Electronics*, vol. 58, no. 10, pp. 4495–4503, 2011.

[109] L. R. d. Araujo, D. R. R. Penido, and F. d. A. Vieira, "A multiphase optimal power flow algorithm for unbalanced distribution systems," *International Journal of Electrical Power & Energy Systems*, vol. 53, pp. 632–642, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0142061513002317

[110] V. M. d. Costa, N. Martins, and J. L. R. Pereira, "Developments in the Newton Raphson power flow formulation based on current injections," *IEEE Transactions on Power Systems*, vol. 14, no. 4, pp. 1320–1326, 1999.

[111] P. A. N. Garcia, J. L. R. Pereira, S. Carneiro *et al.*, "Three-phase power flow calculations using the current injection method," *IEEE Transactions on Power Systems*, vol. 15, no. 2, pp. 508–514, 2000.

[112] D. R. R. Penido, L. R. d. Araujo, S. Carneiro *et al.*, "Three-phase power flow based on four-conductor current injection method for unbalanced distribution networks," *IEEE Transactions on Power Systems*, vol. 23, no. 2, pp. 494–503, 2008.

[113] G. L. Torres and V. H. Quintana, "An interior-point method for nonlinear optimal power flow using voltage rectangular coordinates," *IEEE Transactions on Power Systems*, vol. 13, no. 4, pp. 1211–1218, 1998.

[114] Energy Networks Association, "Active Network Management - Good Practice Guide," ENA, Report, 2015.

[115] F. Olivier, P. Aristidou, D. Ernst, and T. Van Cutsem, "Active management of low-voltage networks for mitigating overvoltages due to photovoltaic units," *IEEE Transactions on Smart Grid*, vol. 7, no. 2, pp. 926–936, March 2016.

[116] T. Boehme, G. P. Harrison, and A. R. Wallace, "Assessment of distribution network limits for non-firm connection of renewable generation," *IET Renewable Power Generation*, vol. 4, no. 1, pp. 64–74, 2010.

[117] S. Jupe, S. Hoda, A. Park *et al.*, "Active management of generation in low-voltage networks," *CIRED - Open Access Proceedings Journal*, vol. 2017, no. 1, pp. 916–919, 2017.

[118] Southern Electric Power Distribution, "Evidence report 250 end point monitors installed sdrc 9.2(a)," SEPD, Report, 2013. [Online]. Available: http://http://www.thamesvalleyvision.co.uk

[119] Western Power Distribution, "OpenLV - SDRC 2.2: Targets networks, market potential and detailed trial design," Western Power Distribution, Report, May 2018.

[120] Lucy GridKey, "Gridkey mcu318 datasheet," 2018. [Online]. Available: https://www.gridkey.co.uk/product/mcu318/

[121] T. C. Green, T. Frost, A. Junyent-Ferré, and P. D. Mitcheson, *Power Electronics for Flexible Distribution Networks*. John Wiley & Sons, Ltd, 2016. [Online]. Available: http://dx.doi.org/10.1002/9781118755471.sgd062

[122] J. M. Bloemink and T. C. Green, "Benefits of distribution-level power electronics for supporting distributed generation growth," *IEEE Transactions on Power Delivery*, vol. 28, no. 2, pp. 911–919, 2013.

[123] R. A. A. d. Graaff, J. M. A. Myrzik, W. L. Kling, and J. H. R. Enslin, "Intelligent nodes in distribution systems - optimizing steady state settings," in *2007 IEEE Lausanne Power Tech*, Conference Proceedings, pp. 391–395.

[124] N. Okada, M. Takasaki, H. Sakai, and S. Katoh, "Development of a 6.6 kV - 1 MVA transformerless loop balance controller," in *2007 IEEE Power Electronics Specialists Conference*, Conference Proceedings, pp. 1087–1091.

[125] E. Romero-Ramos, A. Gomez-Exposito, A. Marano-Marcolini *et al.*, "Assessing the loadability of active distribution networks in the presence of dc controllable links," *IET Generation, Transmission & Distribution*, vol. 5, no. 11, pp. 1105–1113, 2011.

[126] J. M. Bloemink and T. C. Green, "Increasing distributed generation penetration using soft normally-open points," in *IEEE PES General Meeting*, 2010, Conference Proceedings, pp. 1–8.

[127] M. Barragan, J. M. Mauricio, A. Marano *et al.*, "Operational benefits of multiterminal dc-links in active distribution networks," in *2012 IEEE Power and Energy Society General Meeting*, Conference Proceedings, pp. 1–6.

[128] D. Sciano, A. Raza, R. Salcedo *et al.*, "Evaluation of DC links on dense-load urban distribution networks," *IEEE Transactions on Power Delivery*, vol. 31, no. 3, pp. 1317–1326, 2016.

[129] N. Bottrell, P. Lang, and T. Green, "Algorithm for soft open points to solve thermal and voltage constraints in low-voltage distribution networks," *CIRED - Open Access Proceedings Journal*, vol. 2017, no. 1, pp. 1567–1570, 2017.

[130] C. Wang, G. Song, P. Li *et al.*, "Optimal siting and sizing of soft open points in active electrical distribution networks," *Applied Energy*, vol. 189, pp. 301–309, 2017. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0306261916318499

[131] W. Cao, J. Wu, N. Jenkins *et al.*, "Benefits analysis of soft open points for electrical distribution network operation," *Applied Energy*, vol. 165, pp. 36–47, 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0306261915015950

[132] J. Perez-Olvera, T. Green, and A. Junyent-Ferre, "Using multi-terminal DC networks to improve the hosting capacity of distribution networks," in *2018 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, Conference Proceedings, pp. 1–6.

[133] A. Lana, P. Nuutinen, T. Kaipia *et al.*, "Prospects of development of lvdc electricity distribution system energy efficiency," June 2015 2015.

[134] D. Nilsson and A. Sannino, "Efficiency analysis of low- and medium- voltage DC distribution systems," in *IEEE Power Engineering Society General Meeting, 2004.*, Conference Proceedings, pp. 2315–2321 Vol.2.

[135] H. Fudeh and C. M. Ong, "A simple and efficient ac-dc load-flow method for multiterminal dc systems," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-100, no. 11, pp. 4389–4396, 1981.

[136] J. Beerten, S. Cole, and R. Belmans, "A sequential AC/DC power flow algorithm for networks containing Multi-terminal VSC HVDC systems," in *IEEE PES General Meeting*, 2010, Conference Proceedings, pp. 1–7.

[137] J. Beerten, D. V. Hertem, and R. Belmans, "VSC MTDC systems with a distributed DC voltage control - a power flow approach," in *2011 IEEE Trondheim PowerTech*, Conference Proceedings, pp. 1–6.

[138] J. Rimez and R. Belmans, "A combined AC/DC optimal power flow algorithm for meshed AC and DC networks linked by VSC converters," *International Transactions on Electrical Energy Systems*, vol. 25, no. 10, pp. 2024–2035, 2015. [Online]. Available: http://dx.doi.org/10.1002/etep.1943

[139] M. Aragüés-Peñalba, A. E. Alvarez, S. G. Arellano, and O. Gomis-Bellmunt, "Optimal power flow tool for mixed high-voltage alternating current and high-voltage direct current systems for grid integration of large wind power plants," *IET Renewable Power Generation*, vol. 9, no. 8, pp. 876–881, 2015.

[140] G. Daelemans, "VSC HVDC in meshed networks," Thesis, KU Leuven, 2008.

[141] W. Feng, L. A. Tuan, L. B. Tjernberg *et al.*, "A new approach for benefit evaluation of Multiterminal VSC-HVDC using a proposed mixed AC/DC optimal power flow," *IEEE Transactions on Power Delivery*, vol. 29, no. 1, pp. 432–443, 2014.

[142] The MathWorks, Inc, "MATLAB Optimization Toolbox™ User's Guide," 2018, the MathWorks, Natick, MA, USA.

[143] IEEE PES Distribution System Analysis Subcommittee's Distribution Test Feeder WG, "Distribution test feeders," 2017, accessed on 05.08.2017. [Online]. Available: https://site.ieee.org/pes-testfeeders/

[144] K. P. Schneider, B. A. Mather, B. C. Pal *et al.*, "Analytic Considerations and Design Basis for the IEEE Distribution Test Feeders," *IEEE Transactions on Power Systems*, vol. 33, no. 3, pp. 3181–3188, 2018.

[145] J. D. Cross and R. Hartshorn, "My electric avenue: Integrating electric vehicles into the electrical networks," in *6th Hybrid and Electric Vehicles Conference (HEVC 2016)*, Conference Proceedings, pp. 1–6.

[146] M. E. Baran and F. F. Wu, "Network reconfiguration in distribution systems for loss reduction and load balancing," *IEEE Transactions on Power Delivery*, vol. 4, no. 2, pp. 1401–1407, 1989.

[147] R. A. Jabr, R. Singh, and B. C. Pal, "Minimum loss network reconfiguration using mixed-integer convex programming," *IEEE Transactions on Power Systems*, vol. 27, no. 2, pp. 1106–1115, 2012.

[148] Western Power Distribution, "WPD DataPortal2," 2019, [Accessed 01-02-2019]. [Online]. Available: https://dataportal2.westernpower.co.uk/

[149] D. Molzahn, "A survey of distributed optimization algorithms for optimal power flow problems," Report, November 2016 2016.

[150] S. H. Low, "Convex relaxation of optimal power flow part I: Formulations and equivalence," *IEEE Transactions on Control of Network Systems*, vol. 1, no. 1, pp. 15–27, 2014.

[151] W. A. Bukhsh, A. Grothey, K. I. M. McKinnon, and P. A. Trodden, "Local solutions of the optimal power flow problem," *IEEE Transactions on Power Systems*, vol. 28, no. 4, pp. 4780–4788, 2013.

[152] D. K. Molzahn and I. A. Hiskens, "Convex relaxations of optimal power flow problems: An illustrative example," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 63, no. 5, pp. 650–660, 2016.

[153] R. Yang and G. Hug-Glanzmann, "Optimal usage of transmission capacity with FACTS devices in the presence of wind generation: A two-stage approach," in *2012 IEEE Power and Energy Society General Meeting*, Conference Proceedings, pp. 1–7.

[154] E. Sáiz-Marín, E. Lobato, and I. Egido, "Optimal voltage control by wind farms using data mining techniques," *IET Renewable Power Generation*, vol. 8, no. 2, pp. 141–150, 2014.

[155] Q. Peng and S. H. Low, "Distributed algorithm for optimal power flow on a radial network," in *53rd IEEE Conference on Decision and Control*, 2014, Conference Proceedings, pp. 167–172.

[156] P. Šulc, S. Backhaus, and M. Chertkov, "Optimal distributed control of reactive power via the alternating direction method of multipliers," *IEEE Transactions on Energy Conversion*, vol. 29, no. 4, pp. 968–977, 2014.

[157] Y. Han, L. Chen, Z. Wang *et al.*, "Fully distributed optimal power flow for unbalanced distribution networks based on admm," in *2016 IEEE International Conference on Power System Technology (POWERCON)*, 2016, Conference Proceedings, pp. 1–6.

[158] Scottish & Southern Electricity Networks, "Low voltage network modelling and analysis environment - closedown report," SSEN, Report, February 2013.

[159] Electricity North West Limited, "Low voltage network solutions - closedown report," ENWL, Report, June 2014.

[160] A. Al-Wakeel, J. Wu, and N. Jenkins, "State estimation of medium voltage distribution networks using smart meter measurements," *Applied Energy*, vol. 184, pp. 207 – 218, 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0306261916314349

[161] T. Hastie;, R. Tibshirani;, and J. Friedman, *The Elements of Statistical Learning*, 2nd ed. Springer-Verlag New York, 2009.

[162] L. Busoniu, R. Babuska, B. D. Schutter, and D. Ernst, *Reinforcement Learning and Dynamic Programming Using Function Approximators*. CRC Press, 2010.

[163] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5–32, 2001. [Online]. Available: https://doi.org/10.1023/A:1010933404324

[164] R. Berk, *Statistical Learning from a Regression Perspective*. Springer-Verlag New York, 2008.

[165] J. H. Friedman, "Greedy function approximation: A gradient boosting machine," *Ann. Statist.*, vol. 29, no. 5, pp. 1189–1232, 2001. [Online]. Available: https://projecteuclid.org:443/euclid.aos/1013203451

[166] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, http://www.deeplearningbook.org.

[167] H. Borchani, G. Varando, C. Bielza, and P. Larrañaga, "A survey on multi-output regression," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 5, no. 5, pp. 216–233, 2015. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/widm.1157

[168] J. Watt, R. Borhani, and A. K. Katsaggelos, *Machine Learning Refined: Foundations, Algorithms, and Applications*, 2nd ed. Cambridge University Press, 2020.

[169] Y. Bengio, P. Delalleau, and C. Simard, "Decision trees do not generalize to new variations," *Computational Intelligence*, vol. 26, no. 4, 2010.

[170] A. Lenail, "NN-SVG," Accessed: 15-08-2018. [Online]. Available: https://alexlenail.me/NN-SVG/

[171] D. Rumelhart, G. Hinton, and R. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, no. 6, pp. 533–536, 1986.

[172] P. J. Werbos, "Backpropagation through time: what it does and how to do it," *Proceedings of the IEEE*, vol. 78, no. 10, pp. 1550–1560, Oct 1990.

[173] D. P. Kingma and L. Ba, "Adam: A method for stochastic optimization," in *Proceedings of the International Conference on Learning Representations (ICLR)*, San Diego, USA, May 2015.

[174] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, vol. 9, Sardinia, Italy, 13–15 May 2010, pp. 249–256.

[175] A. Navarro-Espinosa and L. Ochoa, "Dissemination document: "Low Voltage Networks Models and Low Carbon Technology Profiles"," The University of Manchester, Electricity North West Limited, Report, June 2015.

[176] V. Mnih, K. Kavukcuoglu, D. Silver *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 02 2015.

[177] O. Vinyals, I. Babuschkin, W. M. Czarnecki *et al.*, "Grandmaster level in StarCraft II using multi-agent reinforcement learning," *Nature*, vol. 575, no. 7782, pp. 350–354, Nov 2019. [Online]. Available: https://doi.org/10.1038/s41586-019-1724-z

[178] M. Fairbank, S. Li, X. Fu *et al.*, "An adaptive recurrent neural-network controller using a stabilization matrix and predictive inputs to solve a tracking problem under disturbances," *Neural Networks*, vol. 49, pp. 74 – 86, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0893608013002438

[179] R. Bellman and R. Kalaba, "On adaptive control processes," *IRE Transactions on Automatic Control*, vol. 4, no. 2, pp. 1–9, 1959.

[180] R. Bellman, "Dynamic programming," *Science*, vol. 153, no. 3731, pp. 34–37, 1966.

[181] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA: MIT Press, 2018.

[182] B. Kim, Y. Zhang, M. van der Schaar, and J. Lee, "Dynamic pricing and energy consumption scheduling with reinforcement learning," *IEEE Transactions on Smart Grid*, vol. 7, no. 5, pp. 2187–2198, Sep. 2016.

[183] F. Ruelens, B. J. Claessens, S. Vandael *et al.*, "Residential demand response of thermostatically controlled loads using batch reinforcement learning," *IEEE Transactions on Smart Grid*, vol. 8, no. 5, pp. 2149–2159, Sep. 2017.

[184] B. V. Mbuwir, F. Spiessens, and G. Deconinck, "Self-learning agent for battery energy management in a residential microgrid," in *2018 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, Oct 2018, pp. 1–6.

[185] E. Mocanu, D. C. Mocanu, P. H. Nguyen *et al.*, "On-line building energy optimization using deep reinforcement learning," *IEEE Transactions on Smart Grid*, vol. 10, no. 4, pp. 3698–3708, July 2019.

[186] D. Ernst, M. Glavic, F. Capitanescu, and L. Wehenkel, "Reinforcement learning versus model predictive control: A comparison on a power system problem," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 39, no. 2, pp. 517–529, 2009.

[187] F. L. Lewis and D. Vrabie, "Reinforcement learning and adaptive dynamic programming for feedback control," *IEEE Circuits and Systems Magazine*, vol. 9, no. 3, pp. 32–50, 2009.

[188] F. Wang, H. Zhang, and D. Liu, "Adaptive dynamic programming: An introduction," *IEEE Computational Intelligence Magazine*, vol. 4, no. 2, pp. 39–47, 2009.

[189] W. B. Powell, *Approximate dynamic programming : solving the curses of dimensionality*. John Wiley & Sons, Ltd, 2011.

[190] D. V. Prokhorov and D. C. Wunsch, "Adaptive critic designs," *IEEE Transactions on Neural Networks*, vol. 8, no. 5, pp. 997–1007, 1997.

[191] P. J. Werbos, "Optimization methods for brain-like intelligent control," in *Proceedings of 1995 34th IEEE Conference on Decision and Control*, vol. 1, Dec 1995, pp. 579–584 vol.1.

[192] P. J. Werbos, "Backwards differentiation in ad and neural nets: Past links and new opportunities," in *Automatic Differentiation: Applications, Theory, and Implementations*, M. Bücker, G. Corliss, U. Naumann *et al.*, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 15–34.

[193] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, p. 436, 2015. [Online]. Available: https://doi.org/10.1038/nature14539

[194] C. J. C. H. Watkins and P. Dayan, "Q-learning," in *Machine Learning*, 1992, pp. 279–292.

[195] Western Power Distribution, "SDRC 2.2 - Targets networks, market potential and detailed trial design," Report, May 2018. [Online]. Available: https://openlv.net/resources/project-information/

[196] K. Arulkumaran, M. P. Deisenroth, M. Brundage, and A. A. Bharath, "Deep reinforcement learning: A brief survey," *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 26–38, 2017.

[197] L. Buşoniu, T. de Bruin, D. Tolić *et al.*, "Reinforcement learning for control: Performance, stability, and deep approximators," *Annual Reviews in Control*, vol. 46, pp. 8–28, 2018.

[198] P. J. Werbos, "Neural networks for control and system identification," in *Proceedings of the 28th IEEE Conference on Decision and Control,*, 1989, pp. 260–265 vol.1.

[199] P. J. Werbos, "Maximizing long-term gas industry profits in two minutes in lotus using neural network methods," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 19, no. 2, pp. 315–333, March 1989.

[200] Y. Lecun and Y. Bengio, *Convolutional networks for images, speech, and time-series*. MIT Press, 1995.

[201] M. Fairbank and E. Alonso, "A comparison of learning speed and ability to cope without exploration between DHP and TD(0)," in *The 2012 International Joint Conference on Neural Networks (IJCNN)*, June 2012, pp. 1–8.

[202] P. Werbos, "Beyond regression: New tools for prediction and analysis in the behavioral sciences," Ph.D. dissertation, Harvard University, Cambridge, MA, Nov. 1974.

[203] P. Werbos, "Approximate dynamic programming for real-time control and neural modeling," in *Handbook of Intelligent Control: Neural Fuzzy and Adaptive Approaches*, D. A. White and D. A. Sofge, Eds. New York: Van Nostrand, 1992, ch. 13.

[204] M. Fairbank, E. Alonso, and D. Prokhorov, "An equivalence between adaptive dynamic programming with a critic and backpropagation through time," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 24, no. 12, pp. 2088–2100, 2013.

[205] B. Widrow, N. K. Gupta, and S. Maitra, "Punish/reward: Learning with a critic in adaptive threshold systems," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. SMC-3, no. 5, pp. 455–465, Sep. 1973.

[206] M. Abadi, A. Agarwal, P. Barham *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," 2015. [Online]. Available: http://download.tensorflow.org/paper/whitepaper2015.pdf

[207] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, "Automatic differentiation in machine learning: a survey," *Journal of Machine Learning Research*, vol. 18, no. 153, pp. 1–43, 2018. [Online]. Available: http://jmlr.org/papers/v18/17-468.html

[208] P. Eberhard and C. Bischof, "Automatic differentiation of numerical integration algorithms," *Mathematics of Computation of the American Mathematical Society*, vol. 68, no. 226, pp. 717–731, 1999.

[209] M. Bartholomew-Biggs, S. Brown, B. Christianson, and L. Dixon, "Automatic differentiation of algorithms," *Journal of Computational and Applied Mathematics*, vol. 124, no. 1, pp. 171 – 190, 2000, numerical Analysis 2000. Vol. IV: Optimization and Nonlinear Equations. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0377042700004222

[210] A. Geron, *Getting started with TensorFlow*, 1st ed. O'Reilly Media, Inc., 2017.

[211] D. Austin, "American Mathematical Society - Feature Column: How to Differentiate with a Computer," Dec 2017. [Online]. Available: http://www.ams.org/publicoutreach/feature-column/fc-2017-12

[212] M. E. Taylor and P. Stone, "Transfer learning for reinforcement learning domains: A survey," *J. Mach. Learn. Res.*, vol. 10, p. 1633–1685, December 2009.

[213] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes *et al.*, Eds. Curran Associates, Inc., 2014, pp. 3320–3328. [Online]. Available: http://papers.nips.cc/paper/5347-how-transferable-are-features-in-deep-neural-networks.pdf

[214] M. Dale, "LV Network Templates For A Low-Carbon Future - Close down report," Report, October 2013 2013. [Online]. Available: https://www.westernpower.co.uk/docs/Innovation/Closed-projects/Network-Templates/LVT-Closedown-Final-2.aspx

[215] M. C. Choy, D. Srinivasan, and R. L. Cheu, "Neural networks for continuous online learning and control," *IEEE Transactions on Neural Networks*, vol. 17, no. 6, pp. 1511–1531, 2006.

[216] A. Nagabandi, C. Finn, and S. Levine, "Deep online learning via meta-learning: Continual adaptation for model-based RL," in *International Conference on Learning Representations*, 2019.

[217] C. Rocha and P. Radatz, "Algorithmo de Fluxo de Potencia do OpenDSS (in Portuguese)," Report, 2017.

# Appendix A

# Fixed-point Algorithm for Power Flow

A numeric example is used in this section to illustrate the procedure followed by the power flow algorithm, from building equivalent circuits for network components to the use of fixed-point iterations. The IEEE 123-bus Test Feeder is used for this, as the data is available in [143].

## A.1 Lines and Cables

Equivalent circuits are used for the overhead lines and the underground cables. They are divided into groups depending on their type. For example, from the data provided in [143], the configurations for the overhead lines and for the underground cables are specified in Table A.1 and in Table A.2, respectively.

Table A.1: Overhead line configurations.

| Configuration | Phasing | Phase Conductor ACSR | Neutral Conductor ACSR | Spacing ID |
|---|---|---|---|---|
| 1 | A B C N | 336,400 26/7 | 4/0 6/1 | 500 |
| 2 | C A B N | 336,400 26/7 | 4/0 6/1 | 500 |
| 3 | B C A N | 336,400 26/7 | 4/0 6/1 | 500 |
| 4 | C B A N | 336,400 26/7 | 4/0 6/1 | 500 |
| 5 | B A C N | 336,400 26/7 | 4/0 6/1 | 500 |
| 6 | A C B N | 336,400 26/7 | 4/0 6/1 | 500 |
| 7 | A C N | 336,400 26/7 | 4/0 6/1 | 505 |
| 8 | A B N | 336,400 26/7 | 4/0 6/1 | 505 |
| 9 | A N | 1/0 | 1/0 | 510 |
| 10 | B N | 1/0 | 1/0 | 510 |
| 11 | C N | 1/0 | 1/0 | 510 |

Table A.2: Underground cables configuration.

| Configuration | Phasing | Cable | Spacing ID |
|---|---|---|---|
| 12 | A B C | 1/0 AA, CN | 515 |

For configuration 1, using the modified Carson's equations, (2.4) and (2.5), the primitive impedance matrix $Z_{prim_1}$ corresponds to (A.1) [39]:

$$\left[Z_{prim_1}\right] = \begin{bmatrix} 0.401 + 1.413j & 0.095 + 0.851j & 0.095 + 0.726j & 0.095 + 0.752j \\ 0.095 + 0.851j & 0.401 + 1.413j & 0.095 + 0.780j & 0.095 + 0.786j \\ 0.095 + 0.726j & 0.095 + 0.780j & 0.401 + 1.413j & 0.095 + 0.767j \\ 0.095 + 0.752j & 0.095 + 0.786j & 0.095 + 0.767j & 0.687 + 1.546j \end{bmatrix} \Omega\text{/mile} \quad \text{(A.1)}$$

After applying the Kron reduction, as in (2.7), the phase impedance matrix $Z_{abc_1}$ is obtained as in (A.2):

$$\left[Z_{abc_1}\right] = \begin{bmatrix} 0.457 + 1.078j & 0.155 + 0.501j & 0.153 + 0.384j \\ 0.155 + 0.501j & 0.466 + 1.048j & 0.158 + 0.423j \\ 0.153 + 0.384j & 0.158 + 0.423j & 0.461 + 1.065j \end{bmatrix} \Omega\text{/mile} \quad \text{(A.2)}$$

Following the equivalent $\pi$ model for lines in (A.3), the primitive admittance matrix $Y_{abc}$ for a 200 ft line can be obtained and is shown in (A.5).

$$\left[Y_{abc}\right] = \begin{bmatrix} [Z_{abc}]^{-1} + \frac{[C_{abc}]}{2} & -[Z_{abc}] \\ -[Z_{abc}] & [Z_{abc}]^{-1} + \frac{[C_{abc}]}{2} \end{bmatrix} \quad \text{(A.3)}$$

The capacitance matrix $C_{abc}$ is the inverse of the potential coefficient matrix, which can be obtained following a similar procedure as above, in this case using the method of conductors and their images. The reader is referred to [39] for more details. The capacitance matrix for configuration 1 is

$$\left[C_{abc_1}\right] = \begin{bmatrix} 0.015 & -0.0049 & -0.0019 \\ -0.0019 & 0.0159 & -0.0031 \\ -0.0019 & -0.0031 & 0.0143 \end{bmatrix} \mu\text{F/mile} \quad \text{(A.4)}$$

The same process is followed to obtain the primitive admittance matrices for the rest of the lines and underground cables. Finally, they are all integrated into the admittance matrix of the system. For the IEEE 123-bus Test Feeder case, the system admittance is a 375x375 matrix and incorporates all the lines and the shunt admittance matrices of the loads and generators.

$$\left[Y_{abc_1}\right] = \begin{bmatrix} 12.84 - 26.35j & -6.23 + 9.04j & -2.25 + 5.90j & -12.84 + 26.63j & 6.23 - 9.14j & 2.25 - 5.9401j \\ -6.23 + 9.04j & 14.36 - 27.30j & -3.82 + 7.19j & 6.23 - 9.14j & -14.36 + 27.60j & 3.82 - 7.2507j \\ -2.25 + 5.90j & -3.82 + 7.19j & 11.54 - 25.43j & 2.25 - 5.94j & 3.82 - 7.25j & -11.54 + 25.70j \\ -12.84 + 26.63j & 6.23 - 9.14j & 2.25 - 5.94j & 12.84 - 26.35j & -6.23 + 9.04j & -2.25 + 5.90j \\ 6.23 - 9.14j & -14.36 + 27.60j & 3.82 - 7.25j & -6.23 + 9.04j & 14.36 - 27.30j & -3.82 + 7.19j \\ 2.25 - 5.94j & 3.82 - 7.25j & -11.54 + 25.70j & -2.25 + 5.90j & -3.82 + 7.19j & 11.54 - 25.43j \end{bmatrix} \text{S} \tag{A.5}$$

## A.2 Loads and Generators

Loads and generators are modelled as Norton equivalents, composed of two parts: 1) a shunt admittance matrix ($Y_{prim,load}$) and 2) a compensation current [95, 217]. This approach is used because it simplifies the incorporation of nonlinear behaviours for the loads and generators (e.g., constant power loads) through the compensation current. Moreover, it facilitates the calculation of power flow for time-series data, where there are relatively small changes in the load/generation between consecutive time-steps [95]. The shunt admittance represents the specified power when the voltage at the load terminals is nominal $V_{nom,load}$, thus it remains fixed during time-series calculations. It can be obtained as in (A.6) and it is included in the shunt admittance matrix of the system.

$$Y_{prim,load} = \frac{(S_i)^*}{V_{nom,load}^2} \tag{A.6}$$

For a single-phase, constant power load with a peak demand of 20 kW and 10 kvar (inductive), and a nominal voltage of 2,401.7 V, the shunt admittance is calculated as in (A.6) and results in (A.7).

$$Y_{prim,load} = \frac{(20,000 + 10,000j)^*}{2,401.7^2} = 0.00347 - 0.00173j \text{ S} \tag{A.7}$$

As described in (2.25), constant power loads consume their nominal active and reactive power regardless of the voltage at their terminals, thus the absorbed current can be obtained using (2.23), which is repeated here for clarity. The voltage from the most recent fixed-point iteration $k$ is used $V_i^k$.

$$I_{inj,load} = \left(\frac{S_i}{V_i^k}\right)^* \tag{A.8}$$

Assuming the voltage at the latest iteration is 0.98 pu with an angle of -0.5 °, the current is obtained as in (A.9).

$$I_{inj,load} = \left(\frac{20,000 + 10,000j}{2,353.6\angle -0.5°}\right)^* = 9.50\angle -27.06° \text{ A} \tag{A.9}$$

Thus, the compensation current of a constant power load can be calculated as in (2.25), repeated here for clarity. The compensation current is obtained as in (A.10), (A.12).

$$I_{comp,load} = [Y_{prim,load}] \times [V_i^k] - I_{inj,load} \tag{A.10}$$

$$I_{comp,load} = (0.00347 - 0.00173j) \times (2,353.6\angle -0.5°) - (9.50\angle -27.06°) \tag{A.11}$$

$$I_{comp,load} = 0.374\angle 151.45° \text{ A} \tag{A.12}$$

A similar process is followed for the rest of the loads and generators. Finally, the compensation currents for all these elements are incorporated into $[I_{nodes}]$ and used for the subsequent fixed-point iteration to update the voltages at all nodes, as in (A.13), repeated here for clarity.

$$[V^k_{nodes}] \leftarrow [Y_{system}]^{-1} \times [I_{nodes}(V^{k-1}_{nodes})] \tag{A.13}$$

## A.3 Fixed-point Algorithm

Given the nonlinear behaviour of the loads and generators, the power flow calculation requires an iterative method to update the compensation currents of the loads and generators at each iteration $k$ with the new voltage estimate. The fixed-point algorithm can be used for this purpose, as summarised in Algorithm 1. The fixed-point algorithm starts by making an initial guess for the voltage vector. The equivalent voltage source at the substation is assumed to have a nominal voltage $V_{base}$ (usually 1.0 pu) and be perfectly balanced (e.g., there is a 120°shift between the phases). For this initial guess the compensation currents of the loads and generators are ignored, however their admittance matrices are included in the system admittance matrix. Thus the voltage vector is calculated as in (A.14), including only the current injected by the voltage source.

$$\begin{bmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \\ \vdots \\ V_n \end{bmatrix} = \begin{bmatrix} Y_{1,1} & Y_{1,2} & \dots Y_{1,n} \\ Y_{2,1} & Y_{2,2} & \dots Y_{2,n} \\ Y_{3,1} & Y_{3,2} & \dots Y_{3,n} \\ Y_{4,1} & Y_{4,2} & \dots Y_{4,n} \\ \vdots & \vdots & \vdots \\ Y_{n,1} & Y_{n,2} & \dots Y_{n,n} \end{bmatrix}^{-1} \begin{bmatrix} I_1 \\ I_2 \\ I_3 \\ I_4 \\ \vdots \\ I_n \end{bmatrix} \tag{A.14}$$

For the IEEE 123-bus Test Feeder case, the equivalent voltage source at the substation is assumed to have a three-phase balanced voltage for the initial iteration, as in (A.15).

$$\left[V_{source}\right] = \begin{bmatrix} 2,401.7\angle 0° \\ 2,401.7\angle -120° \\ 2,401.7\angle 120° \end{bmatrix} \text{ V} \tag{A.15}$$

Similar to the load and generators, the voltage source is also modelled as an equivalent Norton and its admittance matrix is included as part of the system admittance matrix. The current of the voltage source

is obtained using its short circuit impedance, as in (A.17). Because the small magnitude of the short circuit impedance used ($Z_{sc} = 0 + j0.0001$ Ohms), it is expected that the magnitude of the current will be large.

$$I_{source,sc} = \left(\frac{V_{source}}{Z_{sc}}\right) = 24,017\,\text{kA} \tag{A.16}$$

The initial voltage vector is obtained as in (A.17) and (A.18).

$$\begin{bmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \\ \vdots \\ V_n \end{bmatrix} = \begin{bmatrix} Y_{1,1} & Y_{1,2} & \dots Y_{1,n} \\ Y_{2,1} & Y_{2,2} & \dots Y_{2,n} \\ Y_{3,1} & Y_{3,2} & \dots Y_{3,n} \\ Y_{4,1} & Y_{4,2} & \dots Y_{4,n} \\ \vdots & \vdots & \vdots \\ Y_{n,1} & Y_{n,2} & \dots Y_{n,n} \end{bmatrix}^{-1} \begin{bmatrix} 24,017,000\angle -90° \\ 24,017,000\angle 150° \\ 24,017,000\angle 30° \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{A.17}$$

$$\begin{bmatrix} V_1 \\ V_2 \\ V_3 \\ V_4 \\ \vdots \\ V_n \end{bmatrix} = \begin{bmatrix} 2,401.7\angle 0° \\ 2,401.7\angle -120° \\ 2,401.7\angle 120° \\ 2,400.6\angle 0.02° \\ \vdots \\ 2,326.9\angle 123.6° \end{bmatrix} \tag{A.18}$$

With this first guess of the voltage vector, the compensation currents of the loads and generators are updated as described in (A.13). The fixed-point iterations are stopped when the difference between the voltage vector of consecutive iterations is less than the error tolerance, as in (2.28), repeated here for clarity (A.19):

$$V_{error,i} = \frac{\left|\left|V_i^{(k)}\right| - \left|V_i^{(k-1)}\right|\right|}{V_{base}} \tag{A.19}$$

↩